arXiv:1604.07581v2 [cs.DS] 11 Jul 2016

Pattern Matching and Consensus Problems on Weighted Sequences and Profiles*

Tomasz Kociumaka[1], Solon P. Pissis[2], and Jakub Radoszewski[†,1,2]

[1]Institute of Informatics, University of Warsaw, Warsaw, Poland,
[kociumaka,jrad]@mimuw.edu.pl
[2]Department of Informatics, King's College London, London, UK,
solon.pissis@kcl.ac.uk

Abstract

We study pattern matching problems on two major representations of uncertain sequences used in molecular biology: weighted sequences (also known as position weight matrices, PWM) and profiles (i.e., scoring matrices). In the simple version, in which only the pattern or only the text is uncertain, we obtain efficient algorithms with theoretically-provable running times using a variation of the lookahead scoring technique. We also consider a general variant of the pattern matching problems in which both the pattern and the text are uncertain. Central to our solution is a special case where the sequences have equal length, called the consensus problem. We propose algorithms for the consensus problem parameterized by the number of strings that match one of the sequences. As our basic approach, a careful adaptation of the classic meet-in-the-middle algorithm for the knapsack problem is used. On the lower bound side, we prove that our dependence on the parameter is optimal up to lower-order terms conditioned on the optimality of the original algorithm for the knapsack problem.

1 Introduction

We study two well-known representations of uncertain texts: *weighted sequences* and *profiles*. A *weighted sequence* (also known as uncertain sequence or position weight matrix, PWM) for every position and every letter of the alphabet specifies the probability of occurrence of this letter at this position; see Fig. 1 for an example. A weighted sequence represents many different strings, each with the probability of occurrence equal to the product of probabilities of its letters at subsequent positions of the weighted sequence. Usually a threshold $\frac{1}{z}$ is specified, and one considers only strings that match the weighted sequence with probability at least $\frac{1}{z}$. A *scoring matrix* (or a profile) of length m is an $m \times \sigma$ matrix. The *score* of a string of length m is the sum of scores in the scoring matrix of the subsequent letters of the string at the respective positions. A string is said to match a scoring matrix if its matching score is above a specified threshold Z.

$X[1]$	$X[2]$	$X[3]$	$X[4]$
$\pi_1^{(X)}(\mathtt{a}) = 1/2$	$\pi_2^{(X)}(\mathtt{a}) = 1$	$\pi_3^{(X)}(\mathtt{a}) = 3/4$	$\pi_4^{(X)}(\mathtt{a}) = 0$
$\pi_1^{(X)}(\mathtt{b}) = 1/2$	$\pi_2^{(X)}(\mathtt{b}) = 0$	$\pi_3^{(X)}(\mathtt{b}) = 1/4$	$\pi_4^{(X)}(\mathtt{b}) = 1$

Figure 1: A weighted sequence X of length 4 over the alphabet $\Sigma = \{\mathtt{a}, \mathtt{b}\}$

*Work supported by the Polish Ministry of Science and Higher Education under the 'Iuventus Plus' program in 2015-2016 grant no 0392/IP3/2015/73.

†The author is a Newton International Fellow.

Weighted Pattern Matching and Profile Matching First of all, we study the standard variants of pattern matching problems on weighted sequences and profiles, in which only the pattern or the text is an uncertain sequence. In the best-known formulation of the Weighted Pattern Matching problem, we are given a weighted sequence of length n, called a text, a solid (standard) string of length m, called a pattern, both over an alphabet of size σ, and a *threshold probability* $\frac{1}{z}$. We are asked to find all positions in the text where the fragment of length m represents the pattern with probability at least $\frac{1}{z}$. Each such position is called an *occurrence* of the pattern in the text; we also say that the fragment of the text and the pattern *match*. The Weighted Pattern Matching problem can be solved in $\mathcal{O}(\sigma n \log m)$ time via the Fast Fourier Transform [5]. In a more general indexing variant of the problem, considered in [1, 12], one can preprocess a weighted text in $\mathcal{O}(nz^2 \log z)$ time to report all occ occurrences of a given solid pattern of length m in $\mathcal{O}(m + occ)$ time. (A similar indexing data structure, which assumes $z = \mathcal{O}(1)$, was presented in [4].) Very recently, the index construction time was reduced to $\mathcal{O}(nz)$ for constant-sized alphabets [2].

In the classic Profile Matching problem, the pattern is an $m \times \sigma$ profile, the text is a solid string of length n, and our task is to find all positions in the text where the fragment of length m has a score above a specified threshold Z. A naive approach to the Profile Matching problem works in $\mathcal{O}(nm + m\sigma)$ time. A broad spectrum of heuristics improving this algorithm in practice is known; for a survey see [16]. One of the principal techniques, coming in different flavours, is *lookahead scoring* that consists in checking if a partial match could possibly be completed by the following highest scoring letters in the scoring matrix and, if not, pruning the naive search. The Profile Matching problem can also be solved in $\mathcal{O}(\sigma n \log m)$ time via the Fast Fourier Transform [17].

Weighted Consensus and Profile Consensus As our most involved contribution, we study a general variant of pattern matching on weighted sequences and the consensus problems on uncertain sequences, which are closely related to the Multichoice Knapsack problem. In the Weighted Consensus problem, given two weighted sequences of the same length, we are to check if there is a string that matches each of them with probability at least $\frac{1}{z}$. A routine to compare user-entered weighted sequences with existing weighted sequences in the database is used, e.g., in JASPAR, a well-known database of PWMs [19]. In the General Weighted Pattern Matching (GWPM) problem, both the pattern and the text are weighted. In the most common definition of the problem (see [3, 12]), we are to find all fragments of the text that give a positive answer to the Weighted Consensus problem with the pattern. The authors of [3] proposed an algorithm for the GWPM problem based on the weighted prefix table that works in $\mathcal{O}(nz^2 \log z + n\sigma)$ time.

In an analogous way to the Weighted Consensus problem, we define the Profile Consensus problem. Here we are to check for the existence of a string that matches both the scoring matrices above threshold Z. The Profile Consensus problem is actually a special case of the well-known (especially in practice) Multichoice Knapsack problem (also known as the Multiple Choice Knapsack problem). In this problem, we are given n classes $C_1, \ldots, C_n$ of at most λ items each—N items in total—each item c characterized by a value $v(c)$ and a weight $v(c)$. The goal is to select one item from each class so that the sums of values and of weights of the items are below two specified thresholds, V and W. (In the more intuitive formulation of the problem, we require the sum of values to be *above* a specified threshold, but here we consider an equivalent variant in which both parameters are symmetric.) The Multichoice Knapsack problem is widely used in practice, but most research concerns approximation or heuristic solutions; see [14] and references therein. As far as exact solutions are concerned, the classic meet-in-the middle approach by Horowitz and Sahni [11], originally designed for the (binary) Knapsack problem, immediately generalizes to an $\mathcal{O}^*(\lambda^{\lceil \frac{n}{2} \rceil})$-time[1] solution for Multichoice Knapsack.

Several important problems can be expressed as special cases of the Multichoice Knapsack problem using folklore reductions (see [14]). This includes the Subset Sum problem, which for a set of n integers asks whether there is a subset summing up to a given integer Q, and the k-Sum problem which, for $k = \mathcal{O}(1)$ classes of λ integers, asks to choose one element from each class so that the selected integers sum up to zero. These reductions give immediate hardness results for the Multichoice Knapsack problem, and they can be adjusted to yield the same consequences for Profile Consensus. For the Subset Sum problem, as

[1]The $\mathcal{O}^*$ notation suppresses factors polynomial with respect to the instance size (encoded in binary).

shown in [7, 10], the existence of an $\mathcal{O}^*(2^{\varepsilon n})$-time solution for every $\varepsilon > 0$ would violate the Exponential Time Hypothesis (ETH) [13, 15]. Moreover, the $\mathcal{O}^*(2^{n/2})$ running time, achieved in [11], has not been improved yet despite much effort. The 3-Sum conjecture [9] and the more general k-Sum conjecture state that the 3-Sum and k-Sum problems cannot be solved in $\mathcal{O}(\lambda^{2-\varepsilon})$ time and $\mathcal{O}(\lambda^{\lceil \frac{k}{2} \rceil (1-\varepsilon)})$ time, respectively, for any $\varepsilon > 0$.

Our Results As the first result, we show how the lookahead scoring technique combined with a data structure for answering longest common prefix queries in a string can be applied to obtain simple and efficient algorithms for the standard pattern matching problems on uncertain sequences. For a weighted sequence, by R we denote the size of its list representation, and by λ the maximal number of letters with score at least $\frac{1}{z}$ at a single position (thus $\lambda \leq \min(\sigma, z)$). In the Profile Matching problem, we set M as the number of strings that match the scoring matrix with score above Z. In general $M \leq \sigma^m$, however, we may assume that for practical data this number is actually much smaller. We obtain the following running times:

- $\mathcal{O}(m\sigma + n \log M)$ for Profile Matching;
- $\mathcal{O}(R \log^2 \log \lambda + n \log z)$ deterministic and $\mathcal{O}(R + n \log z)$ randomized (Las Vegas, failure with probability R^{-c} for any given constant c) for Weighted Pattern Matching.

The more complex part of our study is related to the consensus problems and to the GWPM problem. Instead of considering Profile Consensus, we study the more general Multichoice Knapsack. We introduce parameters based on the number of solutions with *feasible* weight or value: $A_V = |\{(c_1, \ldots, c_n) : c_i \in C_i \text{ for all } i = 1, \ldots, n, \sum_i v(c_i) \leq V\}|$, that is, the number of choices of one element from each class that satisfy the value threshold; $A_W = |\{(c_1, \ldots, c_n) : c_i \in C_i \text{ for all } i = 1, \ldots, n, \sum_i w(c_i) \leq W\}|$; $A = \max(A_V, A_W)$, and $a = \min(A_V, A_W)$. We obtain algorithms with the following complexities:

- $\mathcal{O}(N + \sqrt{a\lambda} \log A)$ for Multichoice Knapsack;
- $\mathcal{O}(R + \sqrt{z\lambda}(\log \log z + \log \lambda))$ for Weighted Consensus and $\mathcal{O}(n\sqrt{z\lambda}(\log \log z + \log \lambda))$ for General Weighted Pattern Matching.

Since $a \leq A \leq \lambda^n$, our running time for Multichoice Knapsack in the worst case matches (up to lower order terms) the time complexities of the fastest known solutions for both Subset Sum (also binary Knapsack) and 3-Sum. The main novel part of our algorithm for Multichoice Knapsack is an appropriate (yet intuitive) notion of ranks of partial solutions. We also provide a simple reduction from Multichoice Knapsack to Weighted Consensus, which lets us transfer the negative results to the GWPM problem.

- The existence of an $\mathcal{O}^*(z^{\varepsilon})$-time solution for Weighted Consensus for every $\varepsilon > 0$ would violate the Exponential Time Hypothesis.
- For every $\varepsilon > 0$, an $\mathcal{O}^*(z^{0.5-\varepsilon})$-time solution for Weighted Consensus would imply an $\mathcal{O}^*(2^{(0.5-\varepsilon)n})$-time algorithm for Subset Sum.
- For every $\varepsilon > 0$, an $\tilde{\mathcal{O}}(R + z^{0.5}\lambda^{0.5-\varepsilon})$-time[2] solution for Weighted Consensus would imply an $\tilde{\mathcal{O}}(\lambda^{2-\varepsilon})$-time algorithm for 3-Sum.

For the higher-order terms our complexities match the conditional lower bounds; therefore, we put significant effort to keep the lower order terms of the complexities as small as possible.

[2]The $\tilde{\mathcal{O}}$ notation ignores factors polylogarithmic with respect to the input size.

Model of Computations For problems on weighted sequences, we assume the word RAM model with word size $w = \Omega(\log n + \log z)$ and $\sigma = n^{\mathcal{O}(1)}$. We consider the log-probability model of representations of weighted sequences, that is, we assume that probabilities in the weighted sequences and the threshold probability $\frac{1}{z}$ are all of the form $c^{\frac{p}{2^{dw}}}$, where c and d are constants and p is an integer that fits in a constant number of machine words. Additionally, the probability 0 has a special representation. The only operations on probabilities in our algorithms are multiplications and divisions, which can be performed exactly in $\mathcal{O}(1)$ time in this model. Our solutions to the Multichoice Knapsack problem only assume the word RAM model with word size $w = \Omega(\log S + \log a)$, where S is the sum of integers in the input instance; this does not affect the $\mathcal{O}^*$ running time.

Structure of the Paper We start with Preliminaries, where we formally introduce the problems and the main notions used throughout the paper. The following three sections describe our algorithms: in Section 3 for Profile Matching and Weighted Pattern Matching; in Section 4 for Profile Consensus; and in Section 5 for Weighted Consensus and General Weighted Pattern Matching. A tailor-made, yet more efficient algorithm for General Weighted Pattern Matching is presented in Section 6. We conclude with Section 7, where we introduce faster algorithms and matching lower bounds for Multichoice Knapsack and GWPM in the case that λ is large.

2 Preliminaries

Let $\Sigma = \{s_1, s_2, \ldots, s_\sigma\}$ be an alphabet of size σ. A *string* S over Σ is a finite sequence of letters from Σ. We denote the length of S by $|S|$ and, for $1 \leq i \leq |S|$, the i-th letter of S by $S[i]$. By $S[i..j]$ we denote the string $S[i] \ldots S[j]$ called a *factor* of S (if $i > j$, then the factor is an empty string). A factor is called a *prefix* if $i = 1$ and a *suffix* if $j = |S|$. For two strings S and T, we denote their concatenation by $S \cdot T$ (ST in short).

For a string S of length n, by $lcp(i, j)$ we denote the length of the longest common prefix of factors $S[i..n]$ and $S[j..n]$. The following fact specifies a known efficient data structure answering such queries. It consists of the suffix array with its inverse, the LCP table and a data structure for range minimum queries on the LCP table; see [6] for details.

Fact 2.1. *Let S be a string of length n over alphabet of size $\sigma = n^{\mathcal{O}(1)}$. After $\mathcal{O}(n)$-time preprocessing, given indices i and j ($1 \leq i, j \leq n$) one can compute $lcp(i, j)$ in $\mathcal{O}(1)$ time.*

The *Hamming distance* between two strings X and Y of the same length, denoted by $d_H(X, Y)$, is the number of positions where the strings have different letters.

2.1 Profiles

In the Profile Matching problem, we consider a *scoring matrix* (a profile) P of size $m \times \sigma$. For $i \in \{1, \ldots, m\}$ and $j \in \{1, \ldots, \sigma\}$, we denote the integer score of the letter s_j at the position i by $P[i, s_j]$. The *matching score* of a string S of length m with the matrix P is

$$\text{Score}(S, P) = \sum_{i=1}^{m} P[i, S[i]].$$

If $\text{Score}(S, P) \geq Z$ for an integer *threshold* Z, then we say that the string S *matches the matrix P with threshold Z*. We denote the number of strings S that math P with threshold Z by $\text{NumStrings}_Z(P)$.

For a string T and a scoring matrix P, we say that P *occurs in T at position i with threshold Z* if $T[i..i + m - 1]$ matches P with threshold Z. We denote the set of all positions where P occurs in T by $Occ_Z(P, T)$. These notions let us define the Profile Matching problem:

PROFILE MATCHING PROBLEM
Input: A string T of length n, a scoring matrix P of size $m \times \sigma$, and a threshold Z.
Output: The set $Occ_Z(P,T)$.
Parameters: $M = \text{NumStrings}_Z(P)$.

2.2 Weighted Sequences

A *weighted sequence* $X = X[1]\dots X[n]$ of length $|X| = n$ over alphabet $\Sigma = \{s_1, s_2, \dots, s_\sigma\}$ is a sequence of sets of pairs of the form:

$$X[i] = \{(s_j,\ \pi_i^{(X)}(s_j))\ :\ j \in \{1,2,\dots,\sigma\}\}.$$

Here, $\pi_i^{(X)}(s_j)$ is the occurrence probability of the letter s_j at the position $i \in \{1,\dots,n\}$. These values are non-negative and sum up to 1 for a given i.

For all our algorithms, it is sufficient that the probabilities sum up to *at most* 1 for each position. Also, the algorithms sometimes produce auxiliary weighted sequences with sum of probabilities being smaller than 1 on some positions.

We denote the maximum number of letters occurring at a single position of the weighted sequence (with non-zero probability) by λ and the total size of the representation of a weighted sequence by R. The standard representation consists of n lists with up to λ elements each, so $R = \mathcal{O}(n\lambda)$. However, the lists can be shorter in general. Also, if the threshold probability $\frac{1}{z}$ is specified, at each position of a weighted sequence it suffices to store letters with probability at least $\frac{1}{z}$, and clearly there are at most z such letters for each position. This reduction can be performed in linear time, so we shall always assume that $\lambda \le z$.

The *probability of matching* of a string S with a weighted sequence X, $|S| = |X| = m$, is

$$\mathcal{P}(S,X) = \prod_{i=1}^{m} \pi_i^{(X)}(S[i]).$$

We say that a string S *matches* a weighted sequence X with probability at least $\frac{1}{z}$, denoted by $S \approx_{\frac{1}{z}} X$, if $\mathcal{P}(S,X) \ge \frac{1}{z}$. Given a weighted sequence T, by $T[i..j]$ we denote weighted sequence, called a *factor* of T, equal to $T[i]\dots T[j]$ (if $i > j$, then the factor is empty). We then say that a string P *occurs* in T at position i if P matches the factor $T[i..i+m-1]$. We also say that P is a $\frac{1}{z}$-*solid factor* of T at position i (a $\frac{1}{z}$-*solid prefix* if $i = 1$ and a $\frac{1}{z}$-*solid suffix* if $j = n$). We denote the set of all positions where P occurs in T by $Occ_{\frac{1}{z}}(P,T)$.

WEIGHTED PATTERN MATCHING PROBLEM
Input: A string P of length m and a weighted sequence T of length n with at most λ letters at each position and R in total, and a threshold probability $\frac{1}{z}$.
Output: The set $Occ_{\frac{1}{z}}(P,T)$.

3 Profile Matching and Weighted Pattern Matching

In this section we present a solution to the PROFILE MATCHING problem. Afterwards, we show that it can be applied for WEIGHTED PATTERN MATCHING as well.

For a scoring matrix P, the *heavy string* of P, denoted $\mathcal{H}(P)$, is constructed by choosing at each position the heaviest letter, that is, the letter with the maximum score (breaking ties arbitrarily). Intuitively, $\mathcal{H}(P)$ is a string that matches P with the maximum score.

Observation 3.1. *If we have* $\text{Score}(S,P) \ge Z$ *for a string* S *of length* m *and an* $m \times \sigma$ *scoring matrix* P, *then* $d_H(\mathcal{H}(P),S) \le \lfloor \log M \rfloor$ *where* $M = \text{NumStrings}_Z(P)$.

Proof. Let $d = d_H(\mathcal{H}(P), S)$. We can construct 2^d strings of length $|S|$ that match P with a score above Z by taking either of the letters $S[j]$ or $\mathcal{H}(P)[j]$ at each position j such that $S[j] \neq \mathcal{H}(P)[j]$. Hence, $2^d \leq M$, which concludes the proof. □

Our solution for the Profile Matching problem works as follows. We first construct $P' = \mathcal{H}(P)$ and the data structure for finding lcp values between suffixes of P' and T. Let the variable s store the matching score of P'. In the p-th step, we calculate the matching score of $T[p..p+m-1]$ by iterating through subsequent mismatches between P' and $T[p..p+m-1]$ and making adequate updates in the matching score s', which starts at $s' = s$. The mismatches are found using lcp-queries. This process terminates when the score s' drops below Z or when all the mismatches have been found. In the end, we include p in $Occ_Z(P,T)$ if $s' \geq Z$. A pseudocode of this approach is given below for completeness.

Theorem 3.2. Profile Matching *problem can be solved in* $\mathcal{O}(m\sigma + n \log M)$ *time.*

Proof. Let us bound the time complexity of the presented algorithm. The heavy string P' can be computed in $\mathcal{O}(m\sigma)$ time. The data structure for *lcp*-queries in $P'T$ can be constructed in $\mathcal{O}(n+m)$ time by Fact 2.1. Each query for $lcp(P'[i..m], T[j..n])$ can then be answered in constant time by a corresponding *lcp*-query in $P'T$, potentially truncated to the end of P'. Finally, for each position p in the text T we will consider at most $\lfloor \log M \rfloor + 1$ mismatches between P' and T, as afterwards the score s' drops below Z due to Observation 3.1. □

Procedure ProfileMatching(P, T, Z)

```
m := |P|;  n := |T|;  Occ := ∅;
P' := H(P);
Compute the data structure for lcp-queries in P'T;
s := Σ_{j=1}^{m} P[j, P'[j]];
for p := 1 to n − m + 1 do
    s' := s;  i := 1;  j := p;
    while s' ≥ Z and i ≤ m do
        Δ := lcp(P'[i..m], T[j..n]);
        i := i + Δ + 1;  j := j + Δ + 1;
        if i ≤ m + 1 then
            s' := s' + P[i − 1, T[j − 1]] − P[i − 1, P'[i − 1]];
    if s' ≥ Z then insert p to Occ
return Occ;
```

Basically the same approach can be used for Weighted Pattern Matching. In a natural way, we extend the notion of a heavy string to weighted sequences. Now we can restate Observation 3.1 in the language of probabilities instead of scores:

Observation 3.3. *If a string P matches a weighted sequence X of the same length with probability at least $\frac{1}{z}$, then $d_H(\mathcal{H}(X), P) \leq \lfloor \log z \rfloor$.*

Comparing to the solution to Profile Matching, we compute the heavy string of the text instead of the pattern. An auxiliary variable α stores the matching probability between a factor of $\mathcal{H}(T)$ and the corresponding factor of T; it needs to be updated when we move to the next position of the text. In the implementation, we perform the following operations on a weighted sequence:

- computing the probability of a given letter at a given position,
- finding the letter with the maximum probability at a given position.

Procedure WeightedPatternMatching(P, T, $\frac{1}{z}$)

$m := |P|$; $n := |T|$; $Occ := \emptyset$;
$T' := \mathcal{H}(T)$;
Compute the data structure for *lcp*-queries in PT';
$\alpha := \prod_{j=1}^{m} \pi_j^{(T)}(T'[j])$;
for $p := 1$ **to** $n - m + 1$ **do**
  $\alpha' := \alpha$; $i := 1$; $j := p$;
  while $\alpha' \geq \frac{1}{z}$ **and** $i \leq m$ **do**
    $\Delta := lcp(P[i..m], T'[j..n])$;
    $i := i + \Delta + 1$; $j := j + \Delta + 1$;
    if $i \leq m + 1$ **then**
      $\alpha' := \alpha' \cdot \pi_{j-1}^{(T)}(P[i-1]) \,/\, \pi_{j-1}^{(T)}(T'[j-1])$;
  if $\alpha' \geq \frac{1}{z}$ **then** insert p to Occ **if** $p \leq n - m$ **then**
    $\alpha := \alpha \cdot \pi_{p+m}^{(T)}(T'[p+m]) \,/\, \pi_p^{(T)}(T'[p])$;
return Occ;

In the standard list representation, the latter can be performed on a single weighted sequence in $\mathcal{O}(1)$ time after $\mathcal{O}(R)$-time preprocessing. We can perform the former in constant time if, in addition to the list representation, we store the letter probabilities in a dictionary implemented using perfect hashing [8]. This way, we can implement the algorithm in $\mathcal{O}(n \log z + R)$ time w.h.p. Alternatively, a deterministic dictionary [18] can be used to obtain a deterministic solution in $\mathcal{O}(R \log^2 \log \lambda + n \log z)$ time. We arrive at the following result.

Theorem 3.4. Weighted Pattern Matching *can be solved in* $\mathcal{O}(R + n \log z)$ *time with high probability by a Las-Vegas algorithm or in* $\mathcal{O}(R \log^2 \log \lambda + n \log z)$ *time deterministically.*

Remark 3.5. In the same complexity one can solve the GWPM problem with a solid text.

4 Profile Consensus as Multichoice Knapsack

Let us start with a precise statement of the Multichoice Knapsack problem.

Multichoice Knapsack Problem
Input: A set $\mathcal{C}$ of N items partitioned into n disjoint classes C_i, each of size at most λ, two integers $v(c)$ and $w(c)$ for each item $c \in \mathcal{C}$, and two thresholds V and W.
Question: Does there exist a *choice* S (a set $S \subseteq \mathcal{C}$ such that $|S \cap C_i| = 1$ for each i) satisfying both $\sum_{c \in S} v(c) \leq V$ and $\sum_{c \in S} w(c) \leq W$?
Parameters: A_V and A_W: the number of choices S satisfying $\sum_{c \in S} v(c) \leq V$ and $\sum_{c \in S} w(c) \leq W$, respectively, as well as $A = \max(A_V, A_W)$ and $a = \min(A_V, A_W)$.

Indeed we see that the Profile Consensus problem reduces to the Multichoice Knapsack problem. For two $m \times \sigma$ scoring matrices, we construct $n = m$ classes of $\lambda = \sigma$ items each, with values equal to the scores of the letters in the first matrix and weights equal to the scores in the second matrix; both thresholds V and W are equal to Z.

For a fixed instance of Multichoice Knapsack, we say that S is a *partial choice* if $|S \cap C_i| \leq 1$ for each class. The set $D = \{i : |S \cap C_i| = 1\}$ is called its *domain*. For a partial choice S, we define $v(S) = \sum_{c \in S} v(c)$ and $w(S) = \sum_{c \in S} w(c)$.

The classic $\mathcal{O}(2^{n/2})$-time solution to the Knapsack problem [11] partitions $D = \{1, \ldots, n\}$ into two domains D_i of size roughly $n/2$, and for each D_i it generates all partial choices S ordered by $v(S)$. Hence,

it reduces the problem to an instance of MULTICHOICE KNAPSACK with 2 classes. It is solved using the following lemma, proved below for completeness.

Lemma 4.1. *The* MULTICHOICE KNAPSACK *problem can be solved in $\mathcal{O}(N)$ time if $n = 2$ and the elements c of C_1 and C_2 are sorted by $v(c)$.*

Proof. Since the items of C_1 and C_2 are sorted by $v(c)$, a single scan through these items lets us remove all irrelevant elements, that is, elements dominated by other elements in their class. Next, for each $c_1 \in C_1$ we compute $c_2 \in C_2$ such that $v(c_2) \leq V - v(c_1)$ but otherwise $v(c_2)$ is largest possible. As we have removed irrelevant elements from C_2, this item also minimizes $w(c_2)$ among all elements satisfying $v(c_2) \leq V - v(c_1)$. Hence, if there is a feasible solution containing c_1, then $\{c_1, c_2\}$ is feasible. If we process elements c_1 by non-decreasing values $v(c_1)$, the values $v(c_2)$ do not increase, and thus the items c_2 can be computed in $\mathcal{O}(N)$ time in total. □

The same approach generalizes to MULTICHOICE KNAPSACK. The partition is chosen to balance the number of partial choices in each domain, and the worst-case time complexity is $\mathcal{O}(\sqrt{Q\lambda})$, where $Q = \prod_{i=1}^{n} |C_i|$ is the number of choices.

Our aim in this section is to replace Q with the parameter a (which never exceeds Q). The overall running time is going to be $\mathcal{O}(N + \sqrt{a\lambda} \log A)$: an overhead of $\mathcal{O}(\log A)$ appears.

Two challenges arise once we adapt the meet-in-the-middle approach: how to restrict the set of partial choices to be generated so that a feasible solution is not missed, and how to define a partition $D = D_1 \cup D_2$ to balance the number of partial choices generated for D_1 and D_2. A natural idea to deal with the first issue is to consider only partial choices with small values $v(S)$ or $w(S)$. This is close to our actual solution, which is based on the notion of *ranks* of partial choices. Our approach to the second problem is to consider multiple partitions: those of the form $D = \{1, \ldots, j\} \cup \{j + 1, \ldots, n\}$ for $1 \leq j \leq n$. This results in an extra $\mathcal{O}(n)$ factor in the time complexity. However, in Section 4.1 we introduce preprocessing that lets us assume that $n = \mathcal{O}(\frac{\log A}{\log \lambda})$. While dealing with these two issues, some further effort is required to avoid few other extra terms in the running time. In case of our algorithm, this is only $\mathcal{O}(\log \lambda)$, which stems from the fact that we need to keep partial solutions ordered by $v(S)$.

For a partial choice S, we define $\mathrm{rank}_V(S)$ as the number of partial choices S' with the same domain for which $v(S') \leq v(S)$. We symmetrically define $\mathrm{rank}_W(S)$. Ranks are introduced as an analogue of probabilities in weighted sequences. Probabilities are multiplicative, while for ranks we have submultiplicativity:

Fact 4.2. *Assume that $S = S_1 \cup S_2$ is a decomposition of a partial choice S into two disjoint subsets. Then* $\mathrm{rank}_V(S_1)\,\mathrm{rank}_V(S_2) \leq \mathrm{rank}_V(S)$ *(and same for* rank_W*).*

Proof. Let D_1 and D_2 be the domains of S_1 and S_2, respectively. For every partial choices S_1' over D_1 and S_2' over D_2 such that $v(S_1') \leq v(S_1)$ and $v(S_2') \leq v(S_2)$, we have $v(S_1' \cup S_2') = v(S_1') + v(S_2') \leq v(S)$. Hence, $S_1' \cup S_2'$ must be counted while determining $\mathrm{rank}_V(S)$. □

For $0 \leq j \leq n$, let $\mathcal{L}_j$ be the list of partial choices with domain $\{1, \ldots, j\}$ ordered by value $v(S)$, and for $\ell > 0$ let $V_{\mathcal{L}_j}^{(\ell)}$ be the value $v(S)$ of ℓ-th element of $\mathcal{L}_j$ (∞ if $\ell > |\mathcal{L}_j|$). Analogously, for $1 \leq j \leq n + 1$, we define $\mathcal{R}_j$ as the list of partial choices over $\{j, \ldots, n\}$ ordered by $v(S)$, and for $r > 0$, $V_{\mathcal{R}_j}^{(r)}$ as the value of the r-th element of $\mathcal{R}_j$ (∞ if $r > |\mathcal{R}_j|$).

The following two observations yield a decomposition of each choice into a single item and two partial solutions of a small rank. In particular, we do not need to know A_V in order to check if the ranks are sufficiently large.

Lemma 4.3. *Let ℓ and r be positive integers such that $V_{\mathcal{L}_j}^{(\ell)} + V_{\mathcal{R}_{j+1}}^{(r)} > V$ for each $0 \leq j \leq n$. For every choice S with $v(S) \leq V$, there is an index $j \in \{1, \ldots, n\}$ and a decomposition $S = L \cup \{c\} \cup R$ such that $v(L) < V_{\mathcal{L}_{j-1}}^{(\ell)}$, $c \in C_j$, and $v(R) < V_{\mathcal{R}_{j+1}}^{(r)}$.*

Proof. Let $S = \{c_1, \ldots, c_n\}$ with $c_i \in C_i$ and, for $0 \le i \le n$, let $S_i = \{c_1, \ldots, c_i\}$. If $v(S_{n-1}) < V^{(\ell)}_{\mathcal{L}_{n-1}}$, we set $L = S_{n-1}$, $c = c_n$, and $R = \emptyset$, satisfying the claimed conditions.

Otherwise, we define j as the smallest index i such that $v(S_i) \ge V^{(\ell)}_{\mathcal{L}_i}$, and we set $L = S_{j-1}$, $c = c_j$, and $R = S \setminus S_j$. The definition of j implies $v(L) < V^{(\ell)}_{\mathcal{L}_{j-1}}$ and $v(L \cup \{c\}) \ge V^{(\ell)}_{\mathcal{L}_j}$. Moreover, we have $v(L \cup \{c\}) + v(R) = v(S) \le V < V^{(\ell)}_{\mathcal{L}_j} + V^{(r)}_{\mathcal{R}_{j+1}}$, and thus $v(R) < V^{(r)}_{\mathcal{R}_{j+1}}$. □

Fact 4.4. *Let $\ell, r > 0$. If $V^{(\ell)}_{\mathcal{L}_j} + V^{(r)}_{\mathcal{R}_{j+1}} \le V$ for some $j \in \{0, \ldots, n\}$, then $\ell \cdot r \le A_V$.*

Proof. Let L and R be the ℓ-th and r-th entry in $\mathcal{L}_j$ and $\mathcal{R}_{j+1}$, respectively. Note that $v(L \cup R) \le V$ implies $\mathrm{rank}_V(L \cup R) \le A_V$ by definition of A_V. Moreover, $\mathrm{rank}_V(L) \ge \ell$ and $\mathrm{rank}_V(R) \ge r$ (the equalities may be sharp due to draws). Now, Fact 4.2 yields the claimed bound. □

Note that $\mathcal{L}_j$ can be obtained by interleaving $|C_j|$ copies of $\mathcal{L}_{j-1}$, where each copy corresponds to extending the choices from $\mathcal{L}_{j-1}$ with a different item. If we were to construct $\mathcal{L}_j$ having access to the whole $\mathcal{L}_{j-1}$, we could proceed as follows. For each $c \in C_j$, we maintain an *iterator* on $\mathcal{L}_{j-1}$ pointing to the first element S on $\mathcal{L}_{j-1}$ for which $S \cup \{c\}$ has not yet been added to $\mathcal{L}_j$. The associated *value* is $v(S \cup \{c\})$. All iterators initially point at the first element of $\mathcal{L}_{j-1}$. Then the next element to append to $\mathcal{L}_j$ is always $S \cup \{c\}$ corresponding to the iterator with minimum value. Having processed this partial choice, we advance the pointer (or remove it, once it has already scanned the whole $\mathcal{L}_{j-1}$). This process can be implemented using a binary heap H_j as a priority queue, so that initialization requires $\mathcal{O}(|C_j|)$ time and outputting a single element takes $\mathcal{O}(\log |C_j|)$ time.

For all $r \ge 0$, let $\mathcal{L}^{(r)}_j$ be the prefix of $\mathcal{L}_j$ of length $\min(r, |\mathcal{L}_j|)$ and $\mathcal{R}^{(r)}_j$ be the prefix of $\mathcal{R}_j$ of length $\min(r, |\mathcal{R}_j|)$. A technical transformation of the procedure stated above leads to an online algorithm that constructs the prefixes $\mathcal{L}^{(r)}_j$ and $\mathcal{R}^{(r)}_j$. Along with each reported partial choice S, the algorithm also computes $w(S)$.

Lemma 4.5. *After $\mathcal{O}(N)$-time initialization, one can construct $\mathcal{L}^{(i)}_1, \ldots, \mathcal{L}^{(i)}_n$ online for $i = 0, 1, \ldots$, spending $\mathcal{O}(n \log \lambda)$ time per each step. Symmetrically, one can construct $\mathcal{R}^{(i)}_1, \ldots, \mathcal{R}^{(i)}_n$ in the same time complexity.*

Proof. Our online algorithm is going to use the same approach as the offline computation of lists $\mathcal{L}^{(r)}_j$. The order of computations is going to be changed, though.

At each step, for $j = 1$ to n we shall extend lists $\mathcal{L}^{(i-1)}_j$ with a single element (unless the whole $\mathcal{L}_j$ has been already generated) from the top of the heap H_j. Note that this way each iterator in H_j always points to an element that is already in $\mathcal{L}^{(i-1)}_{j-1}$ or to the first element that has not been yet added to $\mathcal{L}_{j-1}$, which is represented by the top of the heap H_{j-1}.

We initialize the heaps as follows: we introduce H_0 which represents the empty choice $\emptyset$ with $v(\emptyset) = 0$. Next, for $j = 1, \ldots, n$ we build the heap H_j representing $|C_j|$ iterators initially pointing to the top of H_{j-1}. The initialization takes $\mathcal{O}(N)$ time in total since a binary heap can be constructed in time linear in its size.

At each step, the lists $\mathcal{L}^{(i-1)}_j$ are extended for consecutive values j from 1 to n. Since $\mathcal{L}^{(i-1)}_{j-1}$ is extended before $\mathcal{L}^{(i-1)}_j$, all iterators in H_j point to the elements of $\mathcal{L}^{(i)}_{j-1}$ while we compute $\mathcal{L}^{(i)}_j$. We take the top of H_j and move it to $\mathcal{L}^{(i)}_j$. Next, we advance the corresponding iterator and update its position in the heap H_j. After this operation, the iterator might point to the top of H_{j-1}. If H_{j-1} is empty, this means that the whole list $\mathcal{L}_{j-1}$ has already been generated and traversed by the iterator. In this case, we remove the iterator.

It is not hard to see that this way we indeed simulate the previous offline solution. A single phase makes $\mathcal{O}(1)$ operations on each heap H_j. The running time is bounded by $\mathcal{O}(\sum_j \log |C_j|) = \mathcal{O}(n \log \lambda)$ at each step of the algorithm. □

The reduction of the following lemma is presented in Section 4.1. Note that we may always assume that $\lambda \le a \le A$. Indeed, if we order the items $c \in C_i$ according to $v(c)$, then only the first A_V of them might belong to a choice S with $v(S) \le V$.

Lemma 4.6. *Given an instance I of the* Multichoice Knapsack *problem, one can compute in $\mathcal{O}(N + \lambda \log A)$ time an equivalent instance I' with $A'_V \le A_V$, $A'_W \le A_W$, $\lambda' \le \lambda$, and $n' = \mathcal{O}(\frac{\log A}{\log \lambda})$.*

Theorem 4.7. Multichoice Knapsack *can be solved in $\mathcal{O}(N + \sqrt{a\lambda}\log A)$ time.*

Proof. Below, we give an algorithm in $\mathcal{O}(N + \sqrt{A_V\lambda}\log A)$ time. The final solution runs it in parallel on the original instance and on the instance with v and V swapped with w and W, waiting until at least one of them terminates.

We increment an integer r starting from 1, maintaining $\ell = \lceil \frac{r}{\lambda} \rceil$ and the lists $\mathcal{L}_j^{(\ell)}$ and $\mathcal{R}_{j+1}^{(r)}$ for $0 \le j \le n$, as long as $V_{\mathcal{L}_j}^{(\ell)} + V_{\mathcal{R}_{j+1}}^{(r)} \le V$ for some j (or until all the lists have been completely generated). By Fact 4.4, we stop at $r = \mathcal{O}(\sqrt{A_V\lambda})$. Lemma 4.6 lets us assume that $n = \mathcal{O}(\frac{\log A}{\log \lambda})$, so the running time of this phase is $\mathcal{O}(N + \sqrt{A_V\lambda}\log A)$ due to Lemma 4.5.

Due to Lemma 4.3, every feasible solution S admits a decomposition $S = L \cup \{c\} \cup R$ with $L \in \mathcal{L}_{j-1}^{(\ell)}$, $c \in C_j$, and $R \in \mathcal{R}_{j+1}^{(r)}$ for some index j. We consider all possibilities for j. For each of them we will reduce searching for S to an instance of the Multichoice Knapsack problem with $N' = \mathcal{O}(\sqrt{A_V\lambda})$ and $n' = 2$. By Lemma 4.1, these instances can be solved in $\mathcal{O}(n\sqrt{A_V\lambda}) = \mathcal{O}(\sqrt{A_V\lambda}\frac{\log A}{\log \lambda})$ time in total.

The items of the j-th instance are going to belong to classes $\mathcal{L}_{j-1}^{(\ell)} \odot C_j$ and $\mathcal{R}_{j+1}^{(r)}$, where $\mathcal{L}_{j-1}^{(\ell)} \odot C_j = \{L \cup \{c\} : L \in \mathcal{L}_{j-1}^{(\ell)}, c \in C_j\}$. The set $\mathcal{L}_{j-1}^{(\ell)} \odot C_j$ can be constructed by merging $|C_j| \le \lambda$ sorted lists each of size $\ell = \mathcal{O}(\sqrt{A_V/\lambda})$, i.e., in $\mathcal{O}(\sqrt{A_V\lambda}\log\lambda)$ time. Summing up over all indices j, this gives $\mathcal{O}(\sqrt{A_V\lambda}\log\lambda\frac{\log A}{\log\lambda}) = \mathcal{O}(\sqrt{A_V\lambda}\log A)$ time.

Clearly, each feasible solution of the constructed instances represents a feasible solution of the initial instance, and by Lemma 4.3, every feasible solution of the initial instance has its counterpart in one of the constructed instances. □

4.1 Proof of Lemma 4.6

Our reduction consists of two steps. Its implementation uses the following notions: For each class C_i, let $v_{\min}(i) = \min\{v(c) : c \in C_i\}$. Also, let $V_{\min} = \sum_{i=1}^{n} v_{\min}(i)$; note that $V_{\min}$ is the smallest possible value $v(S)$ of a choice S. We symmetrically define $w_{\min}(i)$ and $W_{\min}$. First, we make sure that $n = \mathcal{O}(\log A)$.

Lemma 4.8. *Given an instance I of the* Multichoice Knapsack *problem, one can compute in linear time an equivalent instance I' with $N' \le N$, $A'_V \le A_V$, $A'_W \le A_W$, $\lambda' \le \lambda$, and $n' \le 2\log A$.*

Proof. Observe that if some class C_i contains a single item c for which both $v(c) = v_{\min}(i)$ and $w(c) = w_{\min}(i)$, then we can greedily include it in the solution S. Hence, we can remove such a class, setting $V := V - v_{\min}(i)$ and $W := W - w_{\min}(i)$. We execute this reduction rule exhaustively, which clearly takes $\mathcal{O}(N)$ time in total and may only decrease the parameters A_V and A_W. After the reduction, each class contains at least two items.

We shall prove that now we can either find out that $A \ge 2^{n/2}$ or that we are dealing with a NO-instance. To decide which case holds, let us define $\Delta_V(i)$ as the difference between the second smallest value in the multiset $\{v(c) : c \in C_i\}$ and $v_{\min}(i)$. We set Δ_V^{mid} as the sum of the $\lceil \frac{n}{2} \rceil$ smallest values $\Delta_V(i)$ for $1 \le i \le n$; analogously we define Δ_W^{mid}.

Claim. If $V_{\min} + \Delta_V^{\text{mid}} \le V$, then $A_V \ge 2^{n/2}$; if $W_{\min} + \Delta_W^{\text{mid}} \le W$, then $A_W \ge 2^{n/2}$; otherwise, we are dealing with a NO-instance.

Proof. First, assume that $V_{\min} + \Delta_V^{\text{mid}} \le V$. This means that there is a choice S with $v(S) \le V$ containing at least $\frac{n}{2}$ items c such that $\text{rank}_V(c) \ge 2$. Fact 4.2 yields $\text{rank}_V(S) \ge 2^{\lceil n/2 \rceil}$ and consequently $A_V \ge 2^{n/2}$, as claimed. Symmetrically, if $W_{\min} + \Delta_W^{\text{mid}} \le W$, then $A_W \ge 2^{n/2}$.

Now, suppose that there is a feasible solution S. As no class contains a single item minimizing both $v(c)$ and $w(c)$, there are at least $\lceil \frac{n}{2} \rceil$ classes for which S contains an item not minimizing $v(c)$, or at least $\lceil \frac{n}{2} \rceil$

classes for which S contains an item not minimizing $w(c)$. Without loss of generality, we assume that the former holds. Let D be the set of at least $\lceil \frac{n}{2} \rceil$ classes i satisfying the condition. If $c \in C_i$ does not minimize $v(c)$, then $v(c) \geq v_{\min}(i) + \Delta_V(i)$. Consequently, $V \geq v(S) = V_{\min} + \sum_{i \in D} \Delta_V(i)$. However, observe that $\sum_{i \in D} \Delta_V(i) \geq \Delta_V^{\text{mid}}$, so $V \geq V_{\min} + \Delta_V^{\text{mid}}$, as claimed. □

The conditions from the claim can be verified in $\mathcal{O}(N)$ time using a linear-time selection algorithm to compute Δ_V^{mid} and Δ_W^{mid}. If any of the first two conditions holds, we return the instance obtained using our reduction. Otherwise, we output a dummy NO-instance. □

Before we proceed with the second reduction, let us introduce an auxiliary notion. An item $c \in C_j$ is *irrelevant* if there is another item $c' \in C_j$ that *dominates* c, i.e., such that $v(c) > v(c')$ and $w(c) > w(c')$. Removing irrelevant items leads to an equivalent instance of the Multichoice Knapsack problem, and it may only decrease the parameters.

Lemma 4.9. *Consider a class of items in an instance of the* Multichoice Knapsack *problem. In linear time, we can remove some irrelevant items from the class so that the resulting class C satisfies* $\max(\text{rank}_V(c), \text{rank}_W(c)) > \frac{1}{3}|C|$ *for each item $c \in C$.*

Proof. First, note that using a linear-time selection algorithm, we can determine for each item c whether $\text{rank}_V(c) \leq \frac{1}{3}|C|$ and whether $\text{rank}_W(c) \leq \frac{1}{3}|C|$. If there is no item satisfying both conditions, we keep C unaltered. Otherwise, we have an item which dominates at least $|C| - \text{rank}_V(c) - \text{rank}_W(c) \geq \frac{1}{3}|C|$ other items. We scan through all items in C and remove those dominated by c. Next, we repeat the algorithm. The running time of a single phase is clearly linear, and since $|C|$ decreases geometrically, the total running time is also linear. □

A straightforward way to decrease the number of classes is to replace two distinct classes C_i, C_j with their Cartesian product $C_i \times C_j$, assuming that the weight of a pair (c_i, c_j) is the sum of weights of c_i and c_j. This clearly leads to an equivalent instance of the Multichoice Knapsack problem, does not alter the parameters A_V, A_W, and decreases n. On the other hand N and λ may increase; the latter happens only if $|C_i| \cdot |C_j| > \lambda$.

These two reduction rules let us implement our reduction procedure which constitutes the proof of Lemma 4.6.

Proof. First, we apply Lemma 4.8 to make sure that $n \leq 2 \log A$ and $N = \mathcal{O}(\lambda \log A)$. We may now assume that $\lambda \geq 3^6$, as otherwise we already have $n = \mathcal{O}(\frac{\log A}{\log \lambda})$.

Throughout the algorithm, whenever there are distinct classes of size at most $\sqrt{\lambda}$, we shall replace them with their Cartesian product. This may happen only $n - 1$ times, and a single execution takes $\mathcal{O}(\lambda)$ time, so the total running time needed for this part is $\mathcal{O}(\lambda \log A)$.

Furthermore, for every class that we get in the input instance or obtain as a Cartesian product, we apply Lemma 4.9. The total running time spent on this is also $\mathcal{O}(\lambda \log A)$.

Having exhaustively applied these reduction rules, we are guaranteed that $\max(\text{rank}_V(c), \text{rank}_W(c)) > \frac{1}{3}\sqrt{\lambda} \geq \lambda^{\frac{1}{3}}$ for items c from all but one class. Without loss of generality, we assume that the classes satisfying this condition are $C_1, \ldots, C_k$.

Recall that $v_{\min}(i)$ and $w_{\min}(i)$ are defined as minimum values and weights of items in class C_i and that $V_{\min}$ and $W_{\min}$ are their sums over all classes. For $1 \leq i \leq k$, we define $\Delta_V(i)$ as the difference between the $\lceil \lambda^{\frac{1}{3}} \rceil$-th smallest value in the multiset $\{v(c) : c \in C_i\}$ and $v_{\min}(i)$. Next, we define Δ_V^{mid} as the sum of the $\lceil \frac{k}{2} \rceil$ smallest values $\Delta_V(i)$. Symmetrically, we define $\Delta_W(i)$ and Δ_W^{mid}. We shall prove a claim analogous to that in the proof of Lemma 4.8.

Claim. If $V_{\min} + \Delta_V^{\text{mid}} \leq V$, then $A_V \geq \lambda^{\frac{1}{6}k}$; if $W_{\min} + \Delta_W^{\text{mid}} \leq W$, then $A_W \geq \lambda^{\frac{1}{6}k}$; otherwise, we are dealing with a NO-instance.

Proof. First, suppose that $V_{\min} + \Delta_V^{\text{mid}} \leq V$. This means that there is a choice S with $v(S) \leq V$ which contains at least $\frac{k}{2}$ items c with $\text{rank}_V(c) \geq \lambda^{\frac{1}{3}}$. By Fact 4.2, the rank of this choice is at least $\lambda^{\frac{1}{6}k}$, so $A_V \geq \lambda^{\frac{1}{6}k}$, as claimed. The proof of the second case is analogous.

Now, suppose that there is a feasible solution $S = \{c_1, \ldots, c_n\}$. For $1 \leq i \leq k$, we have $\text{rank}_V(c_i) \geq \lambda^{\frac{1}{3}}$ or $\text{rank}_W(c_i) \geq \lambda^{\frac{1}{3}}$. Consequently, $\text{rank}_V(c_i) \geq \lambda^{\frac{1}{3}}$ holds for at least $\lceil \frac{k}{2} \rceil$ classes or $\text{rank}_W(c_i) \geq \lambda^{\frac{1}{3}}$ holds for at least $\lceil \frac{k}{2} \rceil$ classes. Without loss of generality, we assume that the former holds. Let D be the set of (at least $\lceil \frac{k}{2} \rceil$) classes i satisfying the condition. For each $i \in D$, we clearly have $v(c_i) \geq v_{\min}(i) + \Delta_V(i)$, while for each $i \notin D$, we have $v(c_i) \geq v_{\min}(i)$. Consequently, $V \geq v(S) \geq V_{\min} + \sum_{i \in D} \Delta_V(i) \geq V_{\min} + \Delta_V^{\text{mid}}$. Hence, $V \geq V_{\min} + \Delta_V^{\text{mid}}$, which concludes the proof. □

The condition from the claim can be verified using a linear-time selection algorithm: first, we apply it for each class to compute $\Delta_V(i)$ and $\Delta_W(i)$, and then, globally, to determine Δ_V^{mid} and Δ_W^{mid}. If one of the first two conditions hold, we return the instance obtained through the reduction. It satisfies $A \geq \lambda^{\frac{1}{6}k}$, i.e., $n \leq 1 + k \leq 1 + 6\frac{\log A}{\log \lambda}$. Otherwise, we construct a dummy NO-instance. □

5 Weighted Consensus and General Weighted Pattern Matching

The Weighted Consensus problem is formally defined as follows.

> Weighted Consensus Problem
> **Input:** Two weighted sequences X and Y of length n with at most λ letters at each position and R in total, and a threshold probability $\frac{1}{z}$.
> **Output:** A string S such that $S \approx_{\frac{1}{z}} X$ and $S \approx_{\frac{1}{z}} Y$ or NONE if no such string exists.

If two weighted sequences satisfy the consensus, we write $X \approx_{\frac{1}{z}} Y$ and say that X *matches* Y with probability $\frac{1}{z}$. With this definition of a match, we extend the notation and the notation $Occ_{\frac{1}{z}}(P, T)$ to arbitrary weighted sequences.

> General Weighted Pattern Matching (GWPM) Problem
> **Input:** Two weighted sequences P and T of length m and n, respectively, with at most λ letters at each position and R in total, and a threshold probability $\frac{1}{z}$.
> **Output:** The set $Occ_{\frac{1}{z}}(P, T)$.

In the case of the GWPM problem, it is more useful to provide an *oracle* that finds solid factors that correspond to particular occurrences of the pattern. Such an oracle, given $i \in Occ_{\frac{1}{z}}(P, T)$, computes a string that matches both P and $T[i..i + m - 1]$.

We say that a string P is a *maximal $\frac{1}{z}$-solid prefix* of a weighted sequence X if P is a $\frac{1}{z}$-solid prefix of X and no string $P' = Ps$, for $s \in \Sigma$, is a $\frac{1}{z}$-solid prefix of X. Our algorithms rely on the following simple combinatorial observation, originally due to Amir et al. [1].

Fact 5.1 ([1])**.** *A weighted sequence has at most z different maximal $\frac{1}{z}$-solid prefixes.*

The Weighted Consensus problem is actually a special case of Multichoice Knapsack. Namely, given an instance of the former, we can create an instance of the latter with n classes C_i, each containing an item $c_{i,s}$ for every letter s which has non-zero probability at position i in both X and Y. We set $v(c_{i,s}) = -\log \pi_i^{(X)}(s)$ and $w(c_{i,s}) = -\log \pi_i^{(Y)}(s)$ for this item, whereas the thresholds are $V = W = \log z$. It is easy to see that this reduction indeed yields an equivalent instance and that it can be implemented in linear time. By Fact 5.1, we have $A \leq z$ for this instance, so Theorem 4.7 yields the following result:

Corollary 5.2. Weighted Consensus *problem can be solved in $\mathcal{O}(R + \sqrt{z\lambda} \log z)$ time.*

The GWPM problem can be clearly reduced to $n+m-1$ instances of Weighted Consensus. This leads to a naive $\mathcal{O}(nR+n\sqrt{z\lambda}\log z)$-time algorithm. Below, we remove the first term in this complexity. Our solution applies the approach used in Section 3 for Weighted Pattern Matching and uses an observation analogous to Observation 3.3.

Observation 5.3. *If X and Y are weighted sequences that match with threshold $\frac{1}{z}$, then $d_H(\mathcal{H}(X),\mathcal{H}(Y)) \leq 2\lfloor \log z \rfloor$. Moreover there exists a consensus string S such that $S[i]=\mathcal{H}(X)[i]=\mathcal{H}(Y)[i]$ unless $\mathcal{H}(X)[i]\neq\mathcal{H}(Y)[i]$.*

Our algorithm starts by computing $P'=\mathcal{H}(P)$ and $T'=\mathcal{H}(T)$ and the data structure for *lcp*-queries in $P'T'$. We try to match P with every factor $T[p..p+m-1]$ of the text. Following Observation 5.3, we check if $d_H(T'[p..p+m-1],P') \leq 2\lfloor \log z \rfloor$. If not, then we know that no match is possible. Otherwise, let D be the set of positions of mismatches between $T'[p..p+m-1]$ and P'. Assume that we store $\alpha=\prod_{j=1}^{m}\pi^{(T)}_{p+j-1}(T'[p+j-1])$ and $\beta=\prod_{j=1}^{m}\pi^{(P)}_{j}(P'[j])$. Then, in $\mathcal{O}(|D|)$ time we can compute $\alpha'=\prod_{j\notin D}\pi^{(T)}_{p+j-1}(T'[p+j-1])$ and $\beta'=\prod_{j\notin D}\pi^{(P)}_{j}(P'[j])$. Now, we only need to check what happens at the positions in D.

If $D=\emptyset$, then it suffices to check if $\alpha\geq\frac{1}{z}$ and $\beta\geq\frac{1}{z}$. Otherwise, we construct two weighted sequences X and Y by selecting only the positions from D in $T[p..p+m-1]$ and in P. We multiply the probabilities of all letters at the first position in X by α' and in Y by β'. It is clear that $X\approx_{\frac{1}{z}}Y$ if and only if $T[p..p+m-1]\approx_{\frac{1}{z}}P$.

Thus, we reduced the GWPM problem to at most $n-m+1$ instances of the Weighted Consensus problem for strings of length $\mathcal{O}(\log z)$. By Corollary 5.2, solving such an instance takes $\mathcal{O}(\lambda\log z+\sqrt{z\lambda}\log z)=\mathcal{O}(\sqrt{z\lambda}\log z)$ time. Our reduction requires $\mathcal{O}(R\log^2\log\lambda)$ time to preprocess the input (as in Theorem 3.4), but this is dominated by the $\mathcal{O}(n\sqrt{z\lambda}\log z)$ total time of solving the Weighted Consensus instances. If we memorize the solutions to all those instances together with the sets of mismatches D that lead to those instances, then we can also implement the oracle for the GWPM problem with $\mathcal{O}(m)$-time queries. In Section 6, we design a tailor-made solution to replace the generic algorithm for the Multichoice Knapsack problem, which lets us improve the $\log z$ factor to $\log\log z+\log\lambda$.

The following reduction from Multichoice Knapsack to Weighted Consensus immediately yields that any significant improvement in the dependence on z and λ in the running time of our algorithm would lead to breaking long-standing barriers for special cases of Multichoice Knapsack.

Lemma 5.4. *Given an instance I of the* Multichoice Knapsack *problem with n classes of size λ, in linear time one can construct an equivalent instance of the* Weighted Consensus *problem with $z=\mathcal{O}(\prod_{i=1}^{n}|C_i|)$ and sequences of length $\mathcal{O}(n)$ over alphabet of size λ.*

Proof. We construct a pair of weighted sequences X,Y of length n over alphabet $\Sigma=\{1,\ldots,\lambda\}$. Intuitively, choosing letter j at position i will correspond to taking the j-th element of C_i to the solution S, which we denote as $c_{i,j}$. Without loss of generality, we assume that weights and values are non-negative. Otherwise, we may subtract $v_{\min}(i)$ from $v(c_{i,j})$ and $w_{\min}(i)$ from $w(c_{i,j})$ for each item $c_{i,j}$, as well $V_{\min}$ from V and $W_{\min}$ from W.

We set M to the smallest power of two such that $M\geq\max(n,V,W)$. Let $p_i^{(X)}(j)=\log\pi_i^{(X)}(j)$ and $p_i^{(Y)}(j)=\log\pi_i^{(Y)}(j)$ for $j\in\Sigma$. For $j\in\{1,\ldots,|C_i|\}$, we set:

$$p_i^{(X)}(j)=-\frac{\lceil M\log|C_i|\rceil+v(c_{i,j})}{M},\quad p_i^{(Y)}(j)=-\frac{\lceil M\log|C_i|\rceil+w(c_{i,j})}{M}.$$

Clearly, $\sum_{j=1}^{|C_i|}2^{p_i^{(X)}(j)}\leq 1$ and $\sum_{j=1}^{|C_i|}2^{p_i^{(Y)}(j)}\leq 1$. Moreover, we set

$$\log z_X=\frac{1}{M}\left(V+\sum_{i=1}^{n}\lceil M\log|C_i|\rceil\right)\quad\text{and}\quad\log z_Y=\frac{1}{M}\left(W+\sum_{i=1}^{n}\lceil M\log|C_i|\rceil\right).$$

By the choice of M, we have $\max(z_X,z_Y)\leq 2^{\frac{1}{M}(\max(V,W)+n)}\prod_{i=1}^{n}|C_i|\leq 4\prod_{i=1}^{n}|C_i|$.

This way, for a string P of length n, we have

$$\log \mathcal{P}(P,X) = -\frac{1}{M}\left(\sum_{i=1}^{n} \lceil M \log |C_i| \rceil + \sum_{i=1}^{n} v(c_{i,P[i]})\right) \geq -\log z_X \iff \sum_{i=1}^{n} v(c_{i,P[i]}) \leq V$$

and

$$\log \mathcal{P}(P,Y) = -\frac{1}{M}\left(\sum_{i=1}^{n} \lceil M \log |C_i| \rceil + \sum_{i=1}^{n} w(c_{i,P[i]})\right) \geq -\log z_Y \iff \sum_{i=1}^{n} w(c_{i,P[i]}) \leq W.$$

Thus, P is a solution to the constructed instance of the Weighted Consensus problem with two threshold probabilities, $\frac{1}{z_X}$ and $\frac{1}{z_Y}$, if and only if $S = \{c_{i,j} : P[i] = j\}$ is a solution to the underlying instance of the Multichoice Knapsack problem. To have a single threshold $z = \max(z_X, z_Y)$, we append an additional position $n+1$ with symbol 1 only, with $p^{(X)}_{n+1}(1) = 0$ and $p^{(Y)}_{n+1}(1) = \log z_Y - \log z_X$ provided that $z_X \geq z_Y$, and symmetrically otherwise.

If one wants to make sure that the probabilities at each position sum up to exactly one, two further letters can be introduced, one of which gathers the remaining probability in X and has probability 0 in Y, and the other gathers the remaining probability in Y, and has probability 0 in X. □

Theorem 5.5. Weighted Consensus *problem is NP-hard and cannot be solved in:*

1. $\mathcal{O}^*(z^{o(1)})$ *time unless the Exponential Time Hypothesis (ETH) fails;*
2. $\mathcal{O}^*(z^{0.5-\varepsilon})$ *time for some* $\varepsilon > 0$*, unless there is an* $\mathcal{O}^*(2^{(0.5-\varepsilon)n})$*-time algorithm for the* Subset Sum *problem;*
3. $\tilde{\mathcal{O}}(R + z^{0.5}\lambda^{0.5-\varepsilon})$ *time for some* $\varepsilon > 0$ *and for* $n = \mathcal{O}(1)$*, unless there is an* $\mathcal{O}(\lambda^{2(1-\varepsilon)})$*-time algorithm for 3-*Sum.

Proof. We use Lemma 5.4 to derive algorithms for the Multichoice Knapsack problem based on hypothetical solutions for Weighted Consensus. Subset Sum is a special case of Multichoice Knapsack with $\lambda = 2$, i.e., $\prod_i |C_i| = 2^n$. Hence, an $\mathcal{O}^*(z^{o(1)})$-time solution for Weighted Consensus would yield an $\mathcal{O}^*(2^{o(n)})$-time algorithm for Subset Sum, which contradicts ETH by the results of Etscheid et al. [7] and Gurari [10]. Similarly, an $\mathcal{O}^*(z^{0.5-\varepsilon})$-time solution for Weighted Consensus would yield an $\mathcal{O}^*(2^{(0.5-\varepsilon)n})$-time algorithm for Subset Sum. Moreover, k-Sum is a special case of Multichoice Knapsack with $n = k = \mathcal{O}(1)$, i.e., $\prod_i |C_i| = \lambda^k$. Hence, an $\tilde{\mathcal{O}}(R + z^{0.5}\lambda^{0.5-\varepsilon})$-time solution for Weighted Consensus yields an $\mathcal{O}(\lambda + \lambda^{1.5+0.5-\varepsilon}) = \mathcal{O}(\lambda^{2-\varepsilon})$-time algorithm for 3-Sum. □

Nevertheless, it might still be possible to improve the dependence on n in the GWPM problem. For example, one may hope to achieve $\tilde{\mathcal{O}}(nz^{0.5-\varepsilon} + z^{0.5})$ time for $\lambda = \mathcal{O}(1)$.

6 Faster GWPM via Short Dissimilar Weighted Consensus

This section provides a faster solution for the General Weighted Pattern Matching problem. The key ingredient is an improved solution for the following Short Dissimilar Weighted Consensus problem:

Short Dissimilar Weighted Consensus (SDWC) Problem
Input: A threshold probability $\frac{1}{z}$ and two weighted sequences X and Y of length $n \leq 2\lfloor \log z \rfloor$ with at most $\lambda \leq z$ letters at each position and such that $\mathcal{H}(X)$ and $\mathcal{H}(Y)$ are *dissimilar*, i.e., $\mathcal{H}(X)[i] \neq \mathcal{H}(Y)[i]$ for each position i.
Output: A string S such that $S \approx_{\frac{1}{z}} X$ and $S \approx_{\frac{1}{z}} Y$ or NONE if no such string exists.

Note that the instances of the WEIGHTED CONSENSUS problem produced by the reduction of Section 5 are actually instances of the SDWC problem. Our tailor-made solution for the SDWC problem works in $\mathcal{O}(\sqrt{z\lambda}(\log\log z + \log\lambda))$ time. It assumes that the letters at each position of the weighted sequences are sorted according to probabilities (in addition to storing the dictionary of letters and probabilities). This can be achieved in $\mathcal{O}(\lambda\log\lambda)$ time for each position. We have just proved:

Lemma 6.1. *The* GWPM *problem and the computation of its oracle can be reduced in* $\mathcal{O}(n\lambda\log\lambda)$ *time to at most* $n-m+1$ *instances of the* SDWC *problem.*

6.1 Combinatorial Prerequisites

Our improvement upon the algorithm of Theorem 4.7 is based on Fact 5.1, whose analogue does not hold for MULTICHOICE KNAPSACK in general. Technically, instead of the notion of maximal $\frac{1}{z}$-solid prefixes, the algorithm relies on *light* $\frac{1}{z}$*-solid prefixes* defined as follows: We say that a string P of length k is a light $\frac{1}{z}$-solid prefix of a weighted sequence X if $k=0$ or P is a $\frac{1}{z}$-solid prefix of X such that $P[k]\neq\mathcal{H}(X)[k]$. We symmetrically define *light* $\frac{1}{z}$*-solid suffixes* of X. Fact 5.1 lets us bound the number of light solid prefixes.

Fact 6.2. *A weighted sequence has at most* z *different light* $\frac{1}{z}$*-solid prefixes.*

Proof. We show a pair of inverse mappings between the set of maximal $\frac{1}{z}$-solid prefixes of a weighted sequence X and the set of light $\frac{1}{z}$-solid prefixes of X. If P is a maximal $\frac{1}{z}$-solid prefix of X, then we obtain a light $\frac{1}{z}$-solid prefix by removing all trailing letters of P that are heavy letters at the corresponding positions in X. For the inverse mapping, we extend each light $\frac{1}{z}$-solid prefix by heavy letters as long as the prefix is $\frac{1}{z}$-solid. □

We use the notions of light $\frac{1}{z}$-solid prefixes and light $\frac{1}{z}$-solid suffixes to express a result that we will use instead of Lemma 4.3 and Fact 4.4.

Lemma 6.3. *Consider an instance of the* SDWC *problem, and let* $z_\ell, z_r \geq 1$ *be real numbers such that* $z_\ell\cdot z_r \geq z$*. Every consensus string* S *can be decomposed into* $S = L\cdot c\cdot C\cdot R$ *such that the following conditions hold for some* $U,V\in\{X,Y\}$*:*

- L *is a light* $\frac{1}{z_\ell}$*-solid prefix of* U*,*
- c *is a single letter,*
- *all characters of* C *are heavy in* V*,*
- R *is a light* $\frac{1}{z_r}$*-solid suffix of* V*.*

Proof. We set L as the longest proper prefix of S which is a $\frac{1}{z_\ell}$-solid prefix of both X and Y, and we define $k := |L|$. Note that L is a light $\frac{1}{z_\ell}$-solid prefix of X or Y, because $\mathcal{H}(X)$ and $\mathcal{H}(Y)$ are dissimilar. If $k=n-1$, we conclude the proof setting $c=S[n]$ and $C=R$ to empty strings.

Otherwise, we have $\mathcal{P}(S[1..k+1],V[1..k+1]) < \frac{1}{z_\ell}$ for $V=X$ or $V=Y$. Since $\mathcal{P}(S,V)\geq\frac{1}{z}$ and $z_\ell\cdot z_r\geq z$, this implies $\mathcal{P}(S[k+2..n],V[k+2..n])\geq\frac{1}{z_r}$, i.e., that $C\cdot R = S[k+2..n]$ is a $\frac{1}{z_r}$-solid suffix of V. We set C as the longest prefix of $S[k+2..n]$ composed of letters heavy in V. This way R is clearly a light $\frac{1}{z_r}$-solid suffix of V. □

6.2 Computing Light Solid Prefixes

We say that a string P is a common $\frac{1}{z}$-solid prefix (suffix) of X and Y if it is a $\frac{1}{z}$-solid prefix (suffix) of both X and Y. A *standard representation* of a common $\frac{1}{z}$-solid prefix P of length k of X and Y is a triple (P,p_1,p_2) such that p_1 and p_2 are the probabilities $p_1=\mathcal{P}(P,X[1..k])$ and $p_2=\mathcal{P}(P,Y[1..k])$. The string P is written using variable-length encoding so that a letter that occurs at a given position with probability p in X or Y has a representation that consists of $\mathcal{O}(\log\frac{1}{p})$ bits. For every position i, the encoding can be

constructed as follows: we sort letters c according to $\max(\pi_i^{(X)}(c), \pi_i^{(Y)}(c))$ and assign subsequent integer identifiers according to this order. This lets us store a $\frac{1}{z}$-solid factor using $\mathcal{O}(\log z)$ bits: we concatenate the variable-length representations of its letters and we store a bit mask of size $\mathcal{O}(\log z)$ that stores the delimiters between the representations of single letters. An analogous representation can be applied also to common $\frac{1}{z}$-solid suffixes. Our assumptions on the model of computations imply that the standard representation takes constant space. Moreover, constant time is sufficient to extend a common $\frac{1}{z}$-solid prefix by a given letter.

The following observation describes longer light solid prefixes in terms of shorter ones.

Observation 6.4. *Let P be a non-empty light $\frac{1}{z}$-solid prefix of X. If one removes its last letter and then removes all the trailing letters which are heavy at the respective positions in X, then a shorter light $\frac{1}{z}$-solid prefix of X is obtained.*

We build upon Observation 6.4 to derive an efficient algorithm constructing light solid prefixes.

Lemma 6.5. *Let $(X, Y, \frac{1}{z})$ be an instance of the* SDWC *problem and let $z' \leq z$. All common $\frac{1}{z}$-solid prefixes of X and Y being light $\frac{1}{z'}$-solid prefixes of X, sorted first by their length and then by the probabilities in X, can be generated in $\mathcal{O}(z'(\log\log z + \log\lambda) + \log^2 z)$ time.*

Proof. For $k \in \{0, \ldots, n\}$, let $\mathcal{B}_k$ be a list of the requested solid prefixes of length k sorted by their probabilities p_1 in X. Fact 6.2 guarantees that $\sum_{k=0}^{n} |\mathcal{B}_k| \leq z'$.

We compute the lists $\mathcal{B}_k$ for subsequent lengths k. We start with $\mathcal{B}_0$ containing the empty string with its probabilities $p_1 = p_2 = 1$. To compute $\mathcal{B}_k$ for $k > 0$, we use Observation 6.4. We consider all candidates $i = k-1, \ldots, 0$ for the length of the shorter light $\frac{1}{z'}$-solid prefix, and then all letters $s \neq \mathcal{H}(X)[k]$ to put at position k of the new light $\frac{1}{z'}$-solid prefix.

For a given i, we iterate over all elements (P, p_1, p_2) of $\mathcal{B}_i$ ordered by the non-increasing probabilities p_1, and try to extend each of them by the heavy letters in X at positions $i+1, \ldots, k-1$ and by the letter s at position k. We process the letters s ordered by $\pi_k^{(X)}(s)$, ignoring the first one ($\mathcal{H}(X)[k]$) and stopping as soon as we do not get a $\frac{1}{z}$-solid prefix of X.

More precisely, with $X' = \mathcal{H}(X)$, we compute

$$p_1' := p_1 \cdot \prod_{j=i+1}^{k-1} \pi_j^{(X)}(X'[j]) \cdot \pi_k^{(X)}(s) \quad \text{and} \quad p_2' := p_2 \cdot \prod_{j=i+1}^{k-1} \pi_j^{(Y)}(X'[j]) \cdot \pi_k^{(Y)}(s),$$

check if $p_1' \geq \frac{1}{z'}$ and $p_2' \geq \frac{1}{z}$, and, if so, insert $(P \cdot X'[i+1..k-1] \cdot s, p_1', p_2')$ at the beginning of a new list $L_{i,s}$, indexed both by the letter s and by the length i of the shorter light $\frac{1}{z'}$-solid prefix. When we encounter an element (P, p_1, p_2) of $\mathcal{B}_i$ and a letter s for which $p_1' < \frac{1}{z'}$, we proceed to the next element of $\mathcal{B}_i$. If this happens for the heaviest letter $s \neq \mathcal{H}(X)[k]$, we stop considering the current list $\mathcal{B}_i$ and proceed to $\mathcal{B}_{i-1}$. The final step consists in merging all the $k\lambda$ lists $L_{i,s}$ in the order of probabilities in X; the result is $\mathcal{B}_k$.

Let us analyse the time complexity of the k-th step of the algorithm. If an element (P, p_1, p_2) and letter s that we consider satisfy $p_1' \geq \frac{1}{z'}$, this accounts for a new light $\frac{1}{z'}$-solid prefix of X. Hence, in total (over all steps) we consider $\mathcal{O}(z')$ such elements. Note that some of these elements may be discarded due to the condition on p_2'.

For each inspected element (P, p_1, p_2), we also consider at most one letter s for which p_1' is not sufficiently large. If this is not the only letter considered for this element, such candidates can be charged to the previous class. The opposite situation may happen once for each list $\mathcal{B}_i$, which may give $\mathcal{O}(k)$ additional operations in the k-th step, $\mathcal{O}(\log^2 z)$ in total.

Thanks to the order in which the lists are considered, the products of probabilities $\prod_{j=i+1}^{k-1} \pi_j^{(X)}(X'[j])$, $\prod_{j=i+1}^{k-1} \pi_j^{(Y)}(X'[j])$ and factors $X'[i+1..k-1]$ can be stored so that the representation of each subsequent light $\frac{1}{\sqrt{z}}$-solid prefix of length k is computed in $\mathcal{O}(1)$ time. Finally, the merging step in the k-th phase takes $\mathcal{O}(|\mathcal{B}_k| \log(k\lambda)) = \mathcal{O}(|\mathcal{B}_k|(\log\log z + \log\lambda))$ time if a binary heap is used.

The time complexity of the whole algorithm is $\mathcal{O}(\log^2 z + \sum_{k=1}^{n} |\mathcal{B}_k|(\log\log z + \log\lambda))$. By the already mentioned Fact 6.2, this is $\mathcal{O}(\log^2 z + z'(\log\log z + \log\lambda))$. □

6.3 Merge-in-the-Middle Implementation

In this section we apply Lemma 6.3 to solve the SDWC problem. We use Lemma 6.5 to generate all candidates for $L \cdot c$ and R, and we apply a divide-and-conquer procedure to fill this with C. Our procedure works for fixed $U, V \in \{X, Y\}$; the algorithm repeats it for all four choices.

Let $\mathcal{L}_i$ denote a list of all common $\frac{1}{z}$-solid prefixes of X and Y obtained by extending a light $\frac{\sqrt{\lambda}}{\sqrt{z}}$-solid prefix of U of length $i-1$ by a single letter s at position i, and let $\mathcal{R}_i$ denote a list of all common $\frac{1}{z}$-solid suffixes of X and Y of length $n-i+1$ that are light $\frac{1}{\sqrt{z\lambda}}$-solid suffixes of V. We assume that the lists $\mathcal{L}_i$ and $\mathcal{R}_i$ are sorted according to the probabilities in U and V, respectively.

Lemma 6.6. *The lists $\mathcal{L}_i$ and $\mathcal{R}_i$ for $i \in \{1, \ldots, n+1\}$ can be computed in $\mathcal{O}(\sqrt{z\lambda}(\log\log z + \log\lambda))$ time. Their total size is $\mathcal{O}(\sqrt{z\lambda})$.*

Proof. $\mathcal{O}(\sqrt{z\lambda}(\log\log z + \log\lambda))$-time computation of the lists $\mathcal{R}_i$ is directly due to Lemma 6.5. As for the lists $\mathcal{L}_i$, we first compute in $\mathcal{O}(\frac{\sqrt{z}}{\sqrt{\lambda}}(\log\log z + \log\lambda))$ time the lists of all light $\frac{\sqrt{\lambda}}{\sqrt{z}}$-solid prefixes of U, sorted by the lengths of strings and then by the probabilities in U, again using Lemma 6.5. Then for each length $i-1$ and for each letter s at the i-th position, we extend all these prefixes by a single letter. This way we obtain λ lists for a given $i-1$ that can be merged according to the probabilities in U to form the list $\mathcal{L}_i$. Generation of the auxiliary lists takes $\mathcal{O}(\frac{\sqrt{z}}{\sqrt{\lambda}} \cdot \lambda) = \mathcal{O}(\sqrt{z\lambda})$ time in total, and merging them using a binary heap takes $\mathcal{O}(\sqrt{z\lambda}\log\lambda)$ time. This way we obtain an $\mathcal{O}(\sqrt{z\lambda}(\log\log z + \log\lambda))$-time algorithm. □

Let $\mathcal{L}^*_{a,b}$ be a list of common $\frac{1}{z}$-solid prefixes of X and Y of length b obtained by taking a common $\frac{1}{z}$-solid prefix from $\mathcal{L}_i$ for some $i \in \{a, \ldots, b\}$ and extending it by $b-i$ letters that are heavy at the respective positions in V. Similarly, $\mathcal{R}^*_{a,b}$ is a list of common $\frac{1}{z}$-solid suffixes of length $n-a+1$ obtained by taking a common $\frac{1}{z}$-solid suffix from $\mathcal{R}_i$ for some $i \in \{a, \ldots, b\}$ and prepending it by $i-a$ letters that are heavy in V. Again, we assume that each of the lists $\mathcal{L}^*_{a,b}$ and $\mathcal{R}^*_{a,b}$ is sorted according to the probabilities in U and V, respectively.

A *basic interval* is an interval $[a, b]$ represented by its endpoints $1 \le a \le b \le n+1$ such that $2^j \mid a-1$ and $b = \min(n+1, a+2^j-1)$ for some integer j called the *layer* of the interval. For every $j = 0, \ldots, \lceil\log n\rceil$, there are $\Theta(\frac{n}{2^j})$ basic intervals and they are pairwise disjoint.

Example 6.7. For $n = 7$, the basic intervals are $[1,1], \ldots, [8,8], [1,2], [3,4], [5,6], [7,8], [1,4], [5,8], [1,8]$.

Lemma 6.8. *The lists $\mathcal{L}^*_{a,b}$ and $\mathcal{R}^*_{a,b}$ for all basic intervals $[a, b]$ use $\mathcal{O}(\sqrt{z\lambda}\log\log z)$ space and can be constructed in $\mathcal{O}(\sqrt{z\lambda}(\log\log z + \log\lambda))$ time.*

Proof. We compute all the lists $\mathcal{L}^*_{a,b}$ and $\mathcal{R}^*_{a,b}$ for consecutive layers $j = 0, \ldots, \lceil\log n\rceil$ of basic intervals $[a, b]$. For $j = 0$, we have $\mathcal{L}^*_{a,a} = \mathcal{L}_a$ and $\mathcal{R}^*_{a,a} = \mathcal{R}_a$. Suppose that we wish to compute $\mathcal{L}^*_{a,b}$ for $a < b$ at layer j (the computation of $\mathcal{R}^*_{a,b}$ is symmetric). Take $c = a + 2^{j-1} - 1$. Let us iterate through all the elements (P, p_1, p_2) of the list $\mathcal{L}^*_{a,c}$, extend each string P by $\mathcal{H}(V)[c+1..b]$, and multiply the probabilities p_1 and p_2 by

$$\prod_{i=c+1}^{b} \pi_i^{(X)}(\mathcal{H}(V)[i]) \quad \text{and} \quad \prod_{i=c+1}^{b} \pi_i^{(Y)}(\mathcal{H}(V)[i]),$$

respectively. If a common $\frac{1}{z}$-solid prefix is obtained, it is inserted at the end of an auxiliary list L. The resulting list L is merged with $\mathcal{L}^*_{c+1,b}$ according to the probabilities in U; the result is $\mathcal{L}^*_{a,b}$.

Thus, we can compute $\mathcal{L}^*_{a,b}$ in time proportional to the sum of lengths of $\mathcal{L}^*_{a,c}$ and $\mathcal{L}^*_{c+1,b}$. (Note that the necessary products of probabilities can be computed in $\mathcal{O}(n) = \mathcal{O}(\log z)$ total time.) For every $j = 1, \ldots, \lceil\log n\rceil$, the total length of the lists from the j-th layer does not exceed the total length of the lists from the $(j-1)$-th layer. By Lemma 6.6, the lists at the 0-th layer have size $\mathcal{O}(\sqrt{z\lambda})$. The conclusion follows from the fact that $\log n = \mathcal{O}(\log\log z)$. □

Next, we provide an analogue of Lemma 4.1.

Lemma 6.9. *Let L and R be lists containing, for some $k \in \{0, \ldots, n\}$, standard representations of common $\frac{1}{z}$-solid prefixes of length k and common $\frac{1}{z}$-solid suffixes of length $n-k$ of X and Y. If the elements of each list are sorted according to non-decreasing probabilities in X or Y, one can check in $\mathcal{O}(|L| + |R|)$ time whether the concatenation of any $\frac{1}{z}$-solid prefix from L and $\frac{1}{z}$-solid suffix from R yields a string S such that $S \approx_{\frac{1}{z}} X$ and $S \approx_{\frac{1}{z}} Y$.*

Proof. First, we filter out dominated elements of the lists, i.e., elements (P, p_1, p_2) such that there exists another element (P', p_1', p_2') with $p_1' > p_1$ and $p_2' > p_2$. After this operation, we make sure that both lists are sorted with respect to the non-decreasing probabilities in X; this might require reversing the list.

For every element (P, p_1, p_2) of L, we compute the first (leftmost) element (P', p_1', p_2') of R such that $p_1 p_1' \geq \frac{1}{z}$. This element maximizes p_2' among all elements satisfying the latter condition. Hence, it suffices to check if $p_2 p_2' \geq \frac{1}{z}$, and if so, report the result $S = PP'$. As the lists are ordered by p_1 and p_1', respectively, all such elements can be computed in $\mathcal{O}(|L| + |R|)$ total time. □

Finally, we are ready to apply a divide-and-conquer approach to the SDWC problem:

Lemma 6.10. *The* SDWC *problem can be solved in $\mathcal{O}(\sqrt{z\lambda}(\log\log z + \log\lambda))$ time.*

Proof. The algorithm goes along Lemma 6.3, considering all choices of U and V. For each of them, we proceed as follows:

First, we compute the lists $\mathcal{L}_i$, $\mathcal{R}_i$ and $\mathcal{L}^*_{a,b}$, $\mathcal{R}^*_{a,b}$ for all basic intervals. By Lemmas 6.6 and 6.8, this takes $\mathcal{O}(\sqrt{z\lambda}(\log\log z + \log\lambda))$ time. In order to find out if there is a feasible solution, it suffices to attempt joining a common $\frac{1}{z}$-solid prefix from $\mathcal{L}_j$ with a common $\frac{1}{z}$-solid suffix from $\mathcal{R}_k$ for some indices $1 \leq j < k \leq n+1$ by heavy letters of V at positions $j+1, \ldots, k-1$.

We use a recursive routine to find such a pair of indices j, k in a basic interval $[a, b]$ which has positive length and therefore can be decomposed into two basic subintervals $[a, c]$ and $[c+1, b]$. Then either $j \leq c < k$, or both indices j, k belong to the same interval $[a, c]$ or $[c+1, b]$. To check the former case, we apply the algorithm of Lemma 6.9 to $L = \mathcal{L}^*_{a,c}$ and $R = \mathcal{R}^*_{c+1,b}$. The two latter cases are solved by recursive calls for the subintervals.

The recursive routine is called first for the basic interval $[1, n+1]$. The computations performed by the routine for the basic intervals at the j-th level take at most the time proportional to the total size of lists $\mathcal{L}^*_{a,b}$, $\mathcal{R}^*_{a,b}$ at the $(j-1)$-th level. Lemma 6.8 shows that the total size of the lists at all levels is $\mathcal{O}(\sqrt{z\lambda}\log\log z)$. Consequently, the whole procedure works in $\mathcal{O}(\sqrt{z\lambda}(\log\log z + \log\lambda))$ time. □

Lemma 6.10 combined with Lemma 6.1 provides an efficient implementation of the General Weighted Pattern Matching.

Theorem 6.11. *The* GWPM *problem can be solved in $\mathcal{O}(n\sqrt{z\lambda}(\log\log z + \log\lambda))$ time. An oracle for the* GWPM *problem using $\mathcal{O}(n\log z)$ space and supporting queries in $\mathcal{O}(m)$ time can be computed within the same time complexity.*

7 Faster Algorithms for Large λ

In this section we analyse the running times of algorithms for the Multichoice Knapsack problem expressed as $\mathcal{O}(n^{\mathcal{O}(1)} \cdot T(a, \lambda))$ for some function T monotone with respect to both arguments. The algorithm of Theorem 4.7 proves that achieving $T(a, \lambda) = \sqrt{a\lambda}$ is possible. On the other hand, if we assume that Subset Sum does not admit an $\mathcal{O}^*(2^{(0.5-\varepsilon)n})$-time solution, then we immediately get that we cannot have $T(a, 2) = \mathcal{O}(a^{0.5-\varepsilon})$ for any $\varepsilon \geq 0$. Similarly, the 3-Sum conjecture implies that $T(\lambda^3, \lambda) = \mathcal{O}(\lambda^{2-\varepsilon})$ is impossible. While this already refutes the possibility of having $T(a, \lambda) = \mathcal{O}(a^{0.5}\lambda^{0.5-\varepsilon})$ across all arguments (a, λ), such a bound may still hold for some special cases covering an infinite number of arguments. For example, we may potentially achieve $T(a, \lambda) = \mathcal{O}((a\lambda)^{0.5-\varepsilon}) = \mathcal{O}(\lambda^{1.5-\varepsilon})$ for $a = \lambda^2$.

Before we prove that this is indeed possible, let us see the consequences of the hardness of 3-Sum and, in general, $(2k-1)$-Sum. For a non-negative integer k, the $(2k-1)$-Sum conjecture refutes $T(\lambda^{2k-1}, \lambda) =$

$\mathcal{O}(\lambda^{k-\varepsilon})$. By monotonicity of T with respect to the first argument, we conclude that $T(\lambda^c, \lambda) = \mathcal{O}(\lambda^{k-\varepsilon})$ is impossible for $c \geq 2k-1$. On the other hand, monotonicity with respect to the second argument shows that $T(\lambda^c, \lambda) = \mathcal{O}(\lambda^{c\frac{k}{2k-1}-\varepsilon})$ is impossible for $c \leq 2k-1$. The lower bounds following from $(2k-1)$-Sum and $(2k+1)$-Sum turn out to meet at $c = 2k-1+\frac{1}{k+1}$; see Fig. 2.

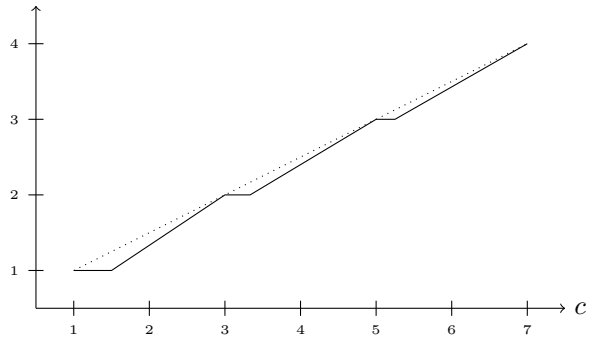

Figure 2: Illustration of the upper bound (dotted) and lower bound (solid) on $\log_\lambda T(\lambda^c, \lambda)$.

Consequently, we have some room between the lower and the upper bound of $\sqrt{a\lambda}$. In the aforementioned case of $a = \lambda^2$, the upper bound is $\lambda^{\frac{3}{2}}$, compared to the lower bound of $\lambda^{\frac{4}{3}-\varepsilon}$. Below, we show that the upper bound can be improved to meet the lower bound. More precisely, we show an algorithm whose running time is $\mathcal{O}(N + (a^{\frac{k+1}{2k+1}} + \lambda^k)\log\lambda \cdot n^k)$ for every positive integer k. Note that $a^{\frac{k+1}{2k+1}} + \lambda^k = \lambda^{c\frac{k+1}{2k+1}} + \lambda^k$, so for $2k-1 \leq c \leq 2k+1$ the running time indeed matches the lower bounds up to the n^k term.

Due to Lemma 4.6, the extra n^k term reduces to $\mathcal{O}((\frac{\log A}{\log\lambda})^k)$, and if we measure the running time using A instead of a, it becomes a constant $(k^{\mathcal{O}(k)})$. In particular, this lets us prove that the GWPM problem can be solved in $\mathcal{O}(n(z^{\frac{k+1}{2k+1}} + \lambda^k)\log\lambda)$ time for any integer $k = \mathcal{O}(1)$, improving upon the solution of Section 5 unless $z = \lambda^{\omega(1)}$ or $z = \lambda^{c\pm o(1)}$ for an odd integer c.

7.1 Algorithm for Multichoice Knapsack

Let us start by discussing the bottleneck of the algorithm of Theorem 4.7 for large λ. The problem is that the size of the classes does not let us partition every choice S into a prefix L and a suffix R with ranks both $\mathcal{O}(\sqrt{A_V})$. Lemma 4.3 leaves us with an extra letter c between L and R, and in the algorithm we append it to the prefix (while generating $\mathcal{L}^{(\ell)}_{j-1} \odot C_j$).

We provide a workaround based on reordering of classes. Our goal is to make sure that items with large rank appear only in a few leftmost classes. For this, we guess the classes of the k items with largest rank (in a feasible solution) and move them to the front. Since this depends on the sought feasible solution, we shall actually verify all $\binom{n}{k}$ possibilities.

Now, our solution considers two cases: For $j > k$, the reordering lets us assume $\mathrm{rank}_V(c) \leq \ell^{\frac{1}{k}}$, so we do not need to consider all items from C_j. For $j \leq k$, on the other hand, we exploit the fact that $|\mathcal{L}^{(\ell)}_{j-1} \odot C_j| \leq \lambda^j$, which at most λ^k.

Combinatorial foundation of this intuition is formalized as a variant of Lemma 4.3:

Lemma 7.1. *Let ℓ and r be positive integers such that $V^{(\ell)}_{\mathcal{L}_j} + V^{(r)}_{\mathcal{R}_{j+1}} > V$ for every $0 \leq j \leq n$. Let $k \in \{1, \ldots, n\}$ and suppose that S is a choice with $v(S) \leq V$ such that $\mathrm{rank}_V(S \cap C_i) \geq \mathrm{rank}_V(S \cap C_j)$ for $i \leq k < j$. There is an index $j \in \{1, \ldots, n\}$ and a decomposition $S = L \cup \{c\} \cup R$ such that $L \in \mathcal{L}^{(\ell)}_{j-1}$, $R \in \mathcal{R}^{(r)}_{j+1}$, $c \in C_j$, and either $\mathrm{rank}_V(c) \leq \ell^{\frac{1}{k}}$ or $j \leq k$.*

Proof. We claim that the decomposition constructed in the proof of Lemma 4.3 satisfies the extra $\mathrm{rank}_V(c) \leq \ell^{\frac{1}{k}}$ condition if $j > k$. Let $S = \{c_1, \ldots, c_n\}$ and $S_i = \{c_1, \ldots, c_i\}$. Obviously $\mathrm{rank}_V(c_i) \geq 1$ for $k < i < j$ and, by the extra assumption, $\mathrm{rank}_V(c_i) \geq \mathrm{rank}_V(c)$ for $1 \leq i \leq k$. Hence, Fact 4.2 yields $\mathrm{rank}_V(S_{j-1}) \geq \mathrm{rank}_V(c)^k$. Simultaneously, we have $v(S_{j-1}) < V^{(\ell)}_{\mathcal{L}_j}$, so $\mathrm{rank}_V(S_{j-1}) < \ell$. Combining these inequalities, we immediately get the claimed bound. □

Theorem 7.2. *For every positive integer $k = \mathcal{O}(1)$, the* Multichoice Knapsack *problem can be solved in $\mathcal{O}(N + (a^{\frac{k+1}{2k+1}} + \lambda^k) \log A (\frac{\log A}{\log \lambda})^k)$ time.*

Proof. As in the proof of Theorem 4.7, we actually provide an algorithm whose running time depends on A_V rather than a. Moreover, Lemma 4.6 lets us assume that $n = \mathcal{O}(\frac{\log A}{\log \lambda})$.

We first guess the k positions where items with largest ranks rank_V are present in the solution S and move these positions to the front. This gives $\binom{n}{k} = \mathcal{O}((\frac{\log A}{\log \lambda})^k)$ possible selections. For each of them, we proceed as follows.

We increment an integer r starting from 1, maintaining $\ell = \lceil r^{\frac{k}{k+1}} \rceil$ and all the lists $\mathcal{L}^{(\ell)}_j$ and $\mathcal{R}^{(r)}_{j+1}$ for $0 \leq j \leq n$, as long as $V^{(\ell)}_{\mathcal{L}_j} + V^{(r)}_{\mathcal{R}_{j+1}} \leq V$ for some j. By Fact 4.4, we stop with $r = \mathcal{O}(A_V^{\frac{k+1}{2k+1}})$ and thus the total time of this phase is $\mathcal{O}(A_V^{\frac{k+1}{2k+1}} \log A)$ due to the online procedure of Lemma 4.5.

By Lemma 7.1, every feasible solution S admits a decomposition $S = L \cup \{c\} \cup R$ for some j; we shall consider all possibilities for j. For each of them, we reduce searching for S to an instance of the Multichoice Knapsack problem with $N' = \mathcal{O}(A_V^{\frac{k+1}{2k+1}} + \lambda^k)$ and $n' = 2$. By Lemma 4.1, these instances can be solved in $\mathcal{O}((A_V^{\frac{k+1}{2k+1}} + \lambda^k) \frac{\log A}{\log \lambda})$ time in total.

For $j \leq k$, the items of the j-th instance are going to belong to classes $\mathcal{L}^{(\ell)}_{j-1} \odot C_j$ and $\mathcal{R}^{(r)}_{j+1}$, where $\mathcal{L}^{(\ell)}_{j-1} \odot C_j = \{L \cup \{c\} : L \in \mathcal{L}^{(\ell)}_{j-1}, c \in C_j\}$. The set $\mathcal{L}^{(\ell)}_{j-1} \odot C_j$ can be sorted by merging $|C_j|$ sorted lists of size at most λ^{j-1} each, i.e., in $\mathcal{O}(\lambda^k \log \lambda)$ time. On the other hand, for $j > k$, we take $\{L \cup \{c\} : L \in \mathcal{L}^{(\ell)}_{j-1}, c \in C_j, \mathrm{rank}_V(c) \leq \ell^{\frac{1}{k}}\}$ and $\mathcal{R}^{(r)}_{j+1}$. The former set can be constructed by merging $\ell^{\frac{1}{k}} = \mathcal{O}(r^{\frac{1}{k+1}})$ sorted lists of size $\mathcal{O}(r^{\frac{k}{k+1}})$ each, i.e., in $\mathcal{O}(r \log \lambda) = \mathcal{O}(A_V^{\frac{k+1}{2k+1}} \log \lambda)$ time.

Summing up over all indices j, this gives $\mathcal{O}((A_V^{\frac{k+1}{2k+1}} + \lambda^k) \log A)$ time for a single selection of the k positions with largest ranks, and $\mathcal{O}((A_V^{\frac{k+1}{2k+1}} + \lambda^k) \log A (\frac{\log A}{\log \lambda})^k)$ in total.

Clearly, each solution of the constructed instances represents a solution of the initial instance, and by Lemma 7.1, every feasible solution of the initial instance has its counterpart in one of the constructed instances.

Before we conclude the proof, let us note that the optimal k does not need to be known in advance. To deal with this issue, we try consecutive integers k and stop the procedure if Fact 4.4 yields that $A_V > \lambda^{2k+1}$, i.e., if r is incremented beyond λ^{k+1}. If the same happens for the other instance of the algorithm (operating on rank_W instead of rank_V), we conclude that $a > \lambda^{2k+1}$, and thus we shall better use larger k. The running time until this point is $\mathcal{O}(\lambda^{k+1} \log \lambda (\frac{\log A}{\log \lambda})^k)$ due to Lemma 4.5. On the other hand, if $r \leq \lambda^{k+1}$, the algorithm behaves as if $a \leq \lambda^{2k+1}$, i.e., runs in $\mathcal{O}(\lambda^{k+1} \log \lambda (\frac{\log A}{\log \lambda})^k)$ time. This workaround (considering all smaller values k) adds extra $\mathcal{O}(\lambda^k \log \lambda (\frac{\log A}{\log \lambda})^{k-1})$ to the time complexity for the optimal value k, which is less than the upper bound on the running time we have for this value k. □

If we are to bound the complexity in terms of A only, the running time becomes

$$\mathcal{O}(N + (A^{\frac{k+1}{2k+1}} + \lambda^k) \log A (\frac{\log A}{\log \lambda})^k).$$

Assumptions that $A \leq \lambda^{2k+1}$ and $k = \mathcal{O}(1)$ let us get rid of the $(\frac{\log A}{\log \lambda})^k$ term, which can be bounded by $(2k+1)^k = \mathcal{O}(1)$. If one of these assumptions is not satisfied, we can improve the bound on the running time anyway: using Theorem 7.2 with increased k if $A > \lambda^{2k+1}$, and using Theorem 4.7 if $k = \omega(1)$.

Corollary 7.3. *Let $k = \mathcal{O}(1)$ be a positive integer such that $A \leq \lambda^{2k+1}$. The* Multichoice Knapsack *problem can be solved in $\mathcal{O}(N + (A^{\frac{k+1}{2k+1}} + \lambda^k) \log \lambda)$ time.*

This leads to the following results for weighted pattern matching:

Theorem 7.4. *Suppose that $\lambda^{2k-1} \leq z \leq \lambda^{2k+1}$ for some positive integer $k = \mathcal{O}(1)$. Then the* SDWC *problem can be solved in $\mathcal{O}((z^{\frac{k+1}{2k+1}} + \lambda^k) \log \lambda)$ time, and the* GWPM *problem can be solved in $\mathcal{O}(n(z^{\frac{k+1}{2k+1}} + \lambda^k) \log \lambda)$ time.*

As we noted at the beginning of this section, Lemma 5.4 implies that any improvement of the dependence of the running time on z or λ by z^ε (equivalently, by λ^ε) would contradict the k-Sum conjecture.

8 Final Remarks

In Section 4, we gave an $\mathcal{O}(N + a^{0.5}\lambda^{0.5} \log A)$-time algorithm for the Multichoice Knapsack problem. Improvement of either exponent to $0.5 - \varepsilon$ would result in a breakthrough for the Subset Sum and 3-Sum problems, respectively. Nevertheless, this does not refute the existence of faster algorithms for some particular values (a, λ) other than those emerging from instances of Subset Sum or 3-Sum. In Section 7, we show an algorithm that is superior if $\frac{\log a}{\log \lambda}$ is a constant other than an odd integer. We also prove it to be optimal (up to lower order terms) for every constant $\frac{\log a}{\log \lambda}$ unless the k-Sum conjecture fails.

References

[1] Amihood Amir, Eran Chencinski, Costas S. Iliopoulos, Tsvi Kopelowitz, and Hui Zhang. Property matching and weighted matching. *Theor. Comput. Sci.*, 395(2-3):298–310, April 2008.

[2] Carl Barton, Tomasz Kociumaka, Jakub Radoszewski, and Solon P. Pissis. Efficient index for weighted sequences, 2016. Accepted to Combinatorial Pattern Matching, CPM 2016.

[3] Carl Barton and Solon P. Pissis. Linear-time computation of prefix table for weighted strings. In Florin Manea and Dirk Nowotka, editors, *Combinatorics on Words, WORDS 2015*, volume 9304 of *LNCS*, pages 73–84. Springer, 2015.

[4] Sudip Biswas, Manish Patil, Sharma V. Thankachan, and Rahul Shah. Probabilistic threshold indexing for uncertain strings. In Evaggelia Pitoura, Sofian Maabout, Georgia Koutrika, Amélie Marian, Letizia Tanca, Ioana Manolescu, and Kostas Stefanidis, editors, *19th International Conference on Extending Database Technology, EDBT 2016*, pages 401–412. OpenProceedings.org, 2016.

[5] Manolis Christodoulakis, Costas S. Iliopoulos, Laurent Mouchard, and Kostas Tsichlas. Pattern matching on weighted sequences. In *Algorithms and Computational Methods for Biochemical and Evolutionary Networks, CompBioNets 2004*, KCL publications, 2004.

[6] Maxime Crochemore, Christophe Hancart, and Thierry Lecroq. *Algorithms on Strings.* Cambridge University Press, New York, NY, USA, 2007.

[7] Michael Etscheid, Stefan Kratsch, Matthias Mnich, and Heiko Röglin. Polynomial kernels for weighted problems. In Giuseppe F. Italiano, Giovanni Pighizzini, and Donald Sannella, editors, *Mathematical Foundations of Computer Science, MFCS 2015, Part II*, volume 9235 of *LNCS*, pages 287–298. Springer, 2015.

[8] Michael L. Fredman, János Komlós, and Endre Szemerédi. Storing a sparse table with $O(1)$ worst case access time. *J. ACM*, 31(3):538–544, 1984.

[9] Anka Gajentaan and Mark H. Overmars. On a class of $O(n^2)$ problems in computational geometry. *Comput. Geom.*, 5:165–185, 1995.

[10] Eitan M. Gurari. *Introduction to the theory of computation.* Computer Science Press, 1989.

[11] Ellis Horowitz and Sartaj Sahni. Computing partitions with applications to the knapsack problem. *J. ACM*, 21(2):277–292, 1974.

[12] Costas S. Iliopoulos, Christos Makris, Yannis Panagis, Katerina Perdikuri, Evangelos Theodoridis, and Athanasios K. Tsakalidis. The weighted suffix tree: An efficient data structure for handling molecular weighted sequences and its applications. *Fundam. Inform.*, 71(2-3):259–277, 2006.

[13] Russell Impagliazzo and Ramamohan Paturi. On the complexity of k-SAT. *J. Comput. Syst. Sci.*, 62(2):367–375, 2001.

[14] Hans Kellerer, Ulrich Pferschy, and David Pisinger. *Knapsack problems.* Springer, 2004.

[15] Daniel Lokshtanov, Dániel Marx, and Saket Saurabh. Lower bounds based on the Exponential Time Hypothesis. *Bulletin of the EATCS*, 105:41–72, 2011.

[16] Cinzia Pizzi and Esko Ukkonen. Fast profile matching algorithms - A survey. *Theor. Comput. Sci.*, 395(2-3):137–157, 2008.

[17] Sanguthevar Rajasekaran, X. Jin, and John L. Spouge. The efficient computation of position-specific match scores with the fast Fourier transform. *J. Comp. Biol.*, 9(1):23–33, 2002.

[18] Milan Ruzic. Constructing efficient dictionaries in close to sorting time. In Luca Aceto, Ivan Damgård, Leslie Ann Goldberg, Magnús M. Halldórsson, Anna Ingólfsdóttir, and Igor Walukiewicz, editors, *Automata, Languages and Programming, ICALP 2008, Part I*, volume 5125 of *LNCS*, pages 84–95. Springer, 2008.

[19] Albin Sandelin, Wynand Alkema, Pär Engström, Wyeth W. Wasserman, and Boris Lenhard. JASPAR: an open-access database for eukaryotic transcription factor binding profiles. *Nucl. Acids Res.*, 32(1):D91–D94, 2004.